

(19)



Europäisches Patentamt  
European Patent Office  
Office européen des brevets



(11)

**EP 0 389 595 B1**

(12)

## EUROPEAN PATENT SPECIFICATION

(45) Date of publication and mention  
of the grant of the patent:

**18.04.2001 Bulletin 2001/16**

(21) Application number: **89910024.2**

(22) Date of filing: **16.08.1989**

(51) Int Cl.7: **G06F 11/22**

(86) International application number:  
**PCT/US89/03528**

(87) International publication number:  
**WO 90/02375 (08.03.1990 Gazette 1990/06)**

(54) **EXPERT SYSTEM FOR DIAGNOSING FAULTS**

EXPERTENSYSTEM FÜR FEHLERDIAGNOSE

SYSTEME EXPERT DE DIAGNOSTIC D'ANOMALIES

(84) Designated Contracting States:  
**DE GB IT NL**

(30) Priority: **19.08.1988 US 234434**

(43) Date of publication of application:  
**03.10.1990 Bulletin 1990/40**

(73) Proprietor: **RAYTHEON COMPANY**  
**Lexington, Massachusetts 02173 (US)**

(72) Inventor: **KIM, Sachol, E.**  
**Yorba Linda, CA 92686 (US)**

(74) Representative: **Witte, Alexander, Dr.-Ing. et al**  
**Witte, Weller, Gahlert, Otten & Steil,**  
**Patentanwälte,**  
**Rotebühlstrasse 121**  
**70178 Stuttgart (DE)**

(56) References cited:

**EP-A- 0 205 873**                      **US-A- 4 649 515**

- **Proceedings of the National Conference on Artificial Intelligence, August 1982, R. DAVIS et al.: "Diagnosis based on Description of Structure and Function", pages 137-142**
- **Expert Systems, Vol. 3, No. 2, April 1986, (Oxford, GB), L.F. PAU: "Survey of Expert Systems for Fault Detection, test Generation, and Maintenance", pages 100-111**
- **Electronic Design, Vol. 34, No. 8, 3 April 1986 (Hasbrouck Heights, New Jersey, US), C. KALME: "A tool Applies the Art of Building Expert Systems to Troubleshooting Boards", pages 159-162, 164,166**
- **Proceedings of the Fourth Conference on Artificial Intelligence Applications, 14-18 March 1988, San Diego (CA, US), CAIA 1988, IEEE, D. LIU: "The HICLASS Software System: a Manufacturing Expert System Shell", pages 256-261**
- **Proceedings of the National Conference on Artificial Intelligence, August 1982, M.R. GENESERETH: "Diagnosis using Hierarchical Design Models", pages 278-283**
- **Proceedings of the Second Conference on Artificial Intelligence Applications, 11-13 December 1985, Miami Beach (Florida, US), CAIA 1985, IEEE, D.L. LARNER: "A Recursive Expert Troubleshooting System Utilizing General and Specific Knowledge", pages 34-41**

Note: Within nine months from the publication of the mention of the grant of the European patent, any person may give notice to the European Patent Office of opposition to the European patent granted. Notice of opposition shall be filed in a written reasoned statement. It shall not be deemed to have been filed until the opposition fee has been paid. (Art. 99(1) European Patent Convention).

**EP 0 389 595 B1**

## Description

**[0001]** The present invention relates to an expert system for diagnosing faults in equipment, said system including a computer having a memory for storing diagnostic routines, a knowledge base accessible by the computer, an inference engine which interfaces with the knowledge base, and a user interface which permits communication between the expert system and an operator,

said knowledge base characterizing the physical, functional and operational relationships of the components and between components of the equipment and information concerning selected operating conditions of the equipment such that said knowledge base contains a conceptual model of the equipment; and

said inference engine interfacing with the operator by way of the user interface and comprising recursive diagnostic routines which operate on the data contained in the knowledge base to generate symptom diagnostic rules in response to operator entered symptoms which are indicative of the faults present in the equipment such that each symptom diagnostic rule correlates a symptom to its fault and is indicative of the fault.

**[0002]** Such an expert system is known from "Proceedings of the National Conference on Artificial Intelligence", August 1982, Pages 278-283; M. R. Genesereth: "Diagnosis Using Hierarchical Design Models".

**[0003]** The present invention relates generally to the field of expert systems, and more particularly to an expert system and method for diagnosing faults in equipment.

**[0004]** Knowledge systems such as expert systems are computer systems that emulate reasoning tasks used by human experts. Such knowledge systems typically use an "inference engine" to interpret encoded knowledge of human experts which is stored in a "knowledge base." If the domain of the knowledge base or scope of the problem is sufficiently narrow, and a sufficiently large body of the knowledge is properly coded in the knowledge base, then the expert system can achieve performance matching or exceeding the ability of a human expert.

**[0005]** Previous attempts to build such expert systems have been based upon a symptom diagnostic rules approach. The essence of this approach is as follows. For every symptom, initial or intermediary, there should correspond an ultimate or intermediary cause. An explicit procedure that leads to a correlation of a symptom to its cause is called a symptom diagnostic rule. These diagnostic rules may be manipulated by a "rule firing" inference engine as follows. Suppose symptom X is exhibited by a given piece of equipment. The inference engine searches for a diagnostic rule in the collection of

such rules which forms the knowledge base. If there is a rule that corresponds to the symptom, then we take the prescribed action Y. If the equipment is restored to normal operation by taking the action Y, the troubleshooting process is completed. If the ultimate cause is not found after the action Y was taken, then a new symptom X<sup>1</sup> is generated. The inference engine then looks for another rule that corresponds to symptom X<sup>1</sup>.

**[0006]** While such expert systems have generally been found useful in the diagnosis of electronic equipment, such systems often had at least one significant disadvantage involving the process for constructing the diagnostic rules. It is almost impossible to anticipate all or even most of the possible symptoms of complex electronic systems and to identify all possible causes for the symptoms. It is also very difficult to develop the correlation procedures between the diagnostic rules. Accordingly, the chief disadvantage of using prior art expert systems to perform diagnosis on electronic equipment is the necessity of constructing the system diagnostic rules.

## SUMMARY OF THE INVENTION

**[0007]** According to the preferred embodiment of the present invention, an expert system is disclosed which diagnoses faults in equipment. The expert system comprises a knowledge base containing information regarding a model of the equipment. Further, the expert system comprises an inference engine which generates a set of symptom diagnostic rules from the information contained in the knowledge base in conjunction with user inputs. The expert system is able to generate information which allows the user of the system to correct faults in the equipment.

**[0008]** Further, "Proceedings of the National Conference on Artificial Intelligence", August 1982, p. 278-283; M.R. Genesereth: "Diagnosis Using Hierarchical Design Models" discloses an expert system for diagnosing faults in computer hardware. The system includes a computer having a memory for storing diagnostic routines, a knowledge base accessible by the computer, an inference engine which interfaces with the knowledge base, and a user interface which permits communication between the expert system and an operator. The knowledge base contains information about the structure of the equipment to be tested. The structure is specified by describing its parts and their interconnections. The structure of each part can in turn be described until one reaches one's "primitive" components which are usually characterized behaviourally. The behaviour of a circuit can usually be expressed in terms of the signal values at the circuit's inputs and outputs. As with structural information, behavioural descriptions are frequently hierarchical, with signals being characterized differently at one level of the structural hierarchy than at another. The knowledge base, thus, contains also formal statements of the relationship between the signals at

these different levels.

**[0009]** Diagnosis is started at the highest level wherein the design model and the symptom input are utilized. If the parts diagnosed to be faulty are replaceable, the diagnosis is complete. Otherwise, the expert system examines those parts at the next lower level of detail. At each level, in a first step, suspects are computed from the design model and the symptom input. If only one suspect is left, the diagnosis is complete. Otherwise, tests are generated wherein a behavioural rule for one of the suspects is utilized so as to generate conclusions using the rules from the devices design model until a proposition of the appropriate form is generated. A proposition relates to the inputs and outputs of a device at one level with their counterparts at the next lower level.

**[0010]** Further, "Expert Systems", Vol. 3, No. 2, April 1986, Oxford, GB, p. 100-111; L.F. Pau: "Survey of Expert Systems for Fault Detection, Test Generation and Maintenance" discloses knowledge-based systems with man-machine interfaces, which use knowledge and reasoning to perform failure detection, testing and maintenance of equipment. The knowledge base contains facts, rules, heuristics and procedural knowledge. An inference engine, being a part of the expert system, comprises reasoning or problem-solving strategies on how to use knowledge to make decisions. Among others, a distributed reasoning expert system architecture is disclosed. This expert system architecture is based on the notion of specialized knowledge bases cooperately solving a specific inference. Each specialized knowledge base may be a production or a structured production system; the corresponding knowledge bases are organized in a hierarchical structure. A blackboard allows all specialist nodes with a structured common area to access information. Such a distributed reasoning architecture comprises a plurality of mini-expert systems and a plurality of intelligent databases interconnected by the blackboard. Particularly, the architecture comprises an intelligent database with a repair plan, another with historical data, another one for spare parts and finally one for TO.

**[0011]** It is the objective of the present invention to provide an improved expert system.

**[0012]** This objective is achieved by the expert system as claimed in Claim 1, wherein said knowledge base is organized in terms of a plurality of files.

**[0013]** According to the preferred embodiment of the present invention, an expert system is disclosed which diagnoses faults in equipment. The expert system comprises a knowledge base containing information regarding a model of the equipment. Further, the expert system comprises an inference engine which generates a set of symptom diagnostic rules from the information contained in the knowledge base in conjunction with user inputs. The expert system is able to generate information which allows the user of the system to correct faults in the equipment.

**[0014]** The advantages of the present invention will

become apparent to one skilled in the art upon reading the following specification with reference to the accompanying drawing, in which:

FIG. 1 is a block diagram of an expert system according to a preferred embodiment of the present invention;

FIG. 2 is a diagram illustrating the connectors representing edges, forming the determinant of the present invention, which connectors are used in FIGS. 3-5;

FIG. 3 is a graphical representation of knowledge elements stored in the knowledge base of the system of FIG. 1;

FIG. 4 is an illustration of a determinant which represents the knowledge about the equipment upon which diagnosis is to be performed;

FIGS. 5a-c are illustrations representing a normalization process which is applied to the determinant shown in FIG. 4;

FIG. 6 is an illustration representing the procedure for collecting and storing information in the knowledge base shown in FIG. 1;

FIG. 7 is an illustration representing an initial state model of the information to be stored in the knowledge base shown in FIG. 1;

FIG. 8 is an illustration representing the transformation of the initial state model shown in FIG. 7 which occurs by fusing out-boundary cells;

FIG. 9 is an illustration representing a tree graph of loops shown in FIG. 8;

FIG. 10 is an illustration representing a Giacombini process which is applied to the tree graph shown in FIG. 9;

FIG. 11 is a tree representation of the determinant of the equipment;

FIG. 12 is an illustration representing a portion of the hierarchical taxonomy of a determinant of the equipment;

FIG. 13 is a flow chart representing the inference engine shown in FIG. 1;

FIG. 14 is a flow chart illustrating the find fault cell routine used by the inference engine;

FIG. 15 is a flow chart representing the find fault routine used by the inference engine;

FIG. 16 is a flow chart representing the blackboard-taxonomy routine used by the inference engine;

FIG. 17 is a flow chart representing the leveler instruction routine used by the inference engine;

FIG. 18 is a flow chart representing the blackboard-knowledge cell routine used by the inference engine;

FIGS. 19a-c is a flow chart which represents the global control heuristic routine used by the inference engine; and

FIGS. 20a-c is a flow chart representing the cell internal control heuristic routine used by the inference engine.

**[0015]** An expert system in accordance with the present invention and generally designated by numeral 24 is shown in FIG. 1. While the expert system 24 may be used for diagnosing faults in electronic equipment, for example, it may also be used for diagnosis in other environments such as mechanical or electromechanical systems, or the like, and hence is not limited to electronic equipment.

**[0016]** The expert system is implemented in a microcomputer which interacts with a user that is preferably familiar with the equipment upon which diagnosis is to be performed. The microcomputer is a general purpose computer and has a central unit including a microprocessor and a random access memory along with one or more floppy disk drives for receiving computer programs and data to be executed and processed.

**[0017]** To permit the user to interact with the expert system, the microcomputer includes a cathode-ray tube display for displaying results and prompting the user for data to be entered via a keyboard. To provide a permanent record for listing of the data when the expert system is in operation, a printer may be provided which is electrically connected to the microcomputer.

**[0018]** The system 24 includes a learning module 36, which is used to gather information from the user regarding the design and operation of the equipment. A user interface 26, which may comprise a keyboard, is provided which allows the user to enter a symptom exhibited by the equipment into the system 24.

**[0019]** The user interface 26 communicates with an inference engine 28, which is a reasoning mechanism that diagnoses malfunctions in the equipment using a determinant or conceptual model of the equipment. As more fully described below, the determinant is used to indicate to the inference engine 28 how the equipment is designed, built and is to be maintained. When the user inputs a given symptom into the system by way of the user interface 26, the inference engine 28 generates an appropriate symptom diagnostic rule which is indicative of the fault. Accordingly, while most expert systems require that symptom diagnostic rules be developed by a knowledge engineer and stored in the knowledge base, the present expert system 24 generates the symptom diagnostic rules.

**[0020]** Given a symptom, the inference engine 28 requests information necessary to perform its diagnosis by way of a blackboarder routine 30. The blackboarder routine 30 is used to access information contained in a knowledge base 32 and composes the information retrieved from the knowledge base 32 on a blackboard 34. To permit the user 12 to locate errors in the knowledge base 32, the system 24 further comprises a justification routine 38. If the system 24 misdiagnoses a particular fault in the equipment, the justification routine 38 permits the user to ask the system 24 to justify the results of its diagnosis. In this way, the user may eliminate errors present in the knowledge base 32.

**[0021]** Based on the information retrieved by the

blackboarder routine 30, the inference engine 28 attempts to isolate the fault. If the fault is isolated, the inference engine 28 indicates to the user that the fault has been isolated and the session is terminated. The user may then use the information obtained from the system 24 to correct the fault in the equipment. If the inference engine 28 is not able to diagnose the fault, it then determines which test should be performed on the equipment and asks the user to perform the test.

**[0022]** If the inference engine 28 does not isolate the fault, it may encounter an infinite loop if it repeatedly examines the same knowledge. The possibility that an infinite loop may occur is minimized, however, since approximately half the fault domain of the equipment is eliminated each time the inference engine 28 conducts a new test. On the other hand, if the inference engine 28 encounters a previously tested location, the system 24 terminates its fault diagnosis, and alerts the user that more information about a certain area of the equipment 24 is needed by generating an ignorance message.

**[0023]** The components of the expert system 24 have been described above in general terms to permit an understanding of the operation thereof. The knowledge base 32 and the inference engine 28 will now be described in greater detail.

## 1. KNOWLEDGE BASE 32

**[0024]** The knowledge base 2 of the expert system 24 will be described from three perspectives: its epistemology; the manner in which information is represented; and the manner in which information is stored.

### 1.1 Epistemology of the Knowledge Base 32

**[0025]** The domain problems in conventional rule-based expert systems are generally the set of all malfunctions that occur in the domain (i.e., the equipment). Each malfunction that occurs in the domain exhibits a set of symptoms so that the domain problem can be resolved if there is a procedure that correlates a given symptom with a particular cause. This type of approach is called a symptom diagnostic rules approach. Accordingly, a domain problem can be resolved if there is a symptom diagnostic rule for every symptom.

**[0026]** Rather than representing knowledge by symptom diagnostic rules, the knowledge base 32 of the present expert system 24 contains two different types of information: that which represents how the equipment is designed, built and is to be maintained; and information indicating whether or not the equipment is operating normally. With these two types of knowledge, the symptom diagnostic rules for most symptoms can be deduced by the inference engine 28. By using this approach, the need for creating large sets of symptom diagnostic rules for storage in the knowledge base 32 is eliminated.

**[0027]** The two types of knowledge described above are contained in sixteen files which comprise the knowl-

edge base 32. Of these sixteen files, two contain objective knowledge relating to information about the physical architecture or anatomy (static view) of the equipment. Objective knowledge consists of specific "devices" and "wires" that link the devices. These two objective knowledge files are contained in a physical device file and a wire list file.

**[0028]** The remaining fourteen files are generally categorized as generated knowledge. Generated knowledge may be classified into three categories. The first category includes information regarding the "dynamic view" of the equipment and consists of a set of specified functions and the relationships among the functions. The dynamic view of the equipment focuses on its functional aspects.

**[0029]** The second type of generated knowledge includes information regarding the "substrates" of the equipment. The equipment may be divided into specific portions which perform a specified function. The portion of the equipment at which the specified function F is performed is termed the substrate F and is denoted by S(F). Accordingly, S(F) correlates the dynamic view of the equipment with its static view.

**[0030]** The third category of generated knowledge includes information covering the normality criteria of the equipment. Once a function F is specified and its substrate S(F) identified, a measurable criteria is needed to allow the system 24 to ascertain whether the substrate S(F) is performing its specified function F correctly. This measurable criteria is referred to as the normality criteria.

**[0031]** The sixteen specific files of the knowledge base 32 are more completely described below.

**[0032]** Taxonomy File: Each defined function which the equipment performs is termed a "knowledge cell" and each knowledge cell has a set of outputs. A downlink (DL) is defined as the carrier of the knowledge cell outputs and is typically a set of wires. The graph formed by the interconnections between knowledge cells and downlinks is called the determinant or system taxonomy of the equipment. A portion of the determinant is a taxonomical piece (TP). Knowledge cells and downlinks must be generated through interpretations of the design documents associated with the equipment as described below. For each taxonomical piece, the taxonomy file contains the name of each knowledge cell, the cell's downlinks, the number of levelers (as defined below) for each composite cell, and the cell type.

**[0033]** Leveler File: A knowledge cell may be decomposable into smaller knowledge cells which reflect a function and its subfunctions. A knowledge cell that can be decomposed is a "composite" knowledge cell, while one which is not decomposable is a "simple" knowledge cell. The leveler file contains the information regarding the decomposition of composite knowledge cells, including the composite cell, the name of the determinant piece, and the name of the root cell of the determinant piece which is the entry point for the composite cell.

**[0034]** Boundary File: Each taxonomical piece contains boundary cells formed from portions of the determinant. The boundary file contains a listing of knowledge cells which are also boundary cells.

**[0035]** Level Jump Conditions (LJC) File: A taxonomical piece is a subset of knowledge cells connected with downlinks and comprises a relatively large number of substrates. The level jump condition file contains a test procedure which the user normally follows to test the substrates collectively (i.e., to determine whether a fault cell exists within the taxonomical piece).

**[0036]** Enablement Conditions File: In order for a substrate S(F) to perform its assigned function F, the substrate S(F) may have a set of conditions that must be satisfied for proper operation. For example, the substrate S(F) may have clock synchronization, voltage or ambient temperature requirements which must be satisfied for the substrate S(F) to operate properly. The enablement conditions file contains information regarding the existence of such conditions for a given knowledge cell.

**[0037]** Maintenance Action File: Even though a substrate S(F) may not contain a fault and all enablement conditions are met, the substrate S(F) still may not perform its assigned function F. The reasons for this may not be deducible from the design documentation alone, in that the problem may be due to the interaction of a complex set of factors which are encountered through experience. For instance, these problems may include intermittent problems, marginal problems, alignment problems or problems caused by abrupt atmospheric or ambient disturbances. The existence of these factors are stored in the maintenance action file. In addition, if there is any known symptom diagnostic rules for the substrate S(F), they are easily incorporated into the knowledge base as maintenance actions in the maintenance action file.

**[0038]** Downlink (DL) File: The downlink file contains information regarding the outputs of each knowledge cell. There are two types of downlinks. One is a set of physical wires which are called the wire links. The other is called a logical link, whose outputs may be probed or inspected on a substrate other than the currently probed substrate.

**[0039]** Test Message (TM) File: For each taxonomical piece which does not contain faulty knowledge cells, a test message is generated which indicates the test procedure that is to be performed on that piece. The test message file is used to store outputs of all the knowledge cells in the taxonomical piece.

**[0040]** Normal Operational State Values (NOSV) File: The normal outputs along the wire links of a knowledge cell with respect to a given test message are called normal operational state values (NOSV) of the knowledge cell. The values are stored in the normal operational state values file.

**[0041]** Logical State Values File: While the NOSV file contains normal operational state values for wire links,

the logical state values file contains normal operational state values of the logical links, such as conditions of surfaces determined through visual inspection, oily surfaces, and displayed information resulting from test message inputs.

**[0042]** Special Probe/Inspection Procedure File: The special probe/inspection procedure file is used to probe or inspect normal operating state values for each knowledge cell, including special equipment and procedures.

**[0043]** Determinant Directory File: As more fully described below, knowledge is represented in the knowledge base in the form of a determinant model. For convenience, the determinant may be cut into many pieces called determinant pieces (DP). The determinant directory is used to inform the inference engine 28 how to proceed from one determinant piece to another. Each determinant piece is uniquely defined by its name, the name of the taxonomical piece with which it is associated, and the pointers or keys to the portions of the level jump conditions file and the test message file associated with the particular taxonomical piece.

**[0044]** GLCH Override File: The GLCH override file contains information that overrides the diagnostic routine performed by the inference engine 28 on a particular knowledge cell to another knowledge cell which is more likely to contain the fault based on knowledge engineering experience and inputs. The GLCH override file may be used to indicate that the inference engine should examine one particular knowledge cell if another particular knowledge cell is encountered.

**[0045]** CICH Override File: The CICH override file contains information that overrides the reasoning process of the inference engine 28 as it troubleshoots within a given knowledge cell. The CICH override file allows the user to instruct the inference engine 28 to look for a particular fault when a certain knowledge cell is determined to be faulty.

**[0046]** Physical Device File: The physical device file contains information regarding the architecture of the equipment (i.e., the devices forming the equipment), including the names of the components which correspond to the hardware nodes of the tree graph structure, including card names, display, frame attach devices such as track balls, and memory, for example.

**[0047]** Wire List File: The wire list file contains information regarding the wires which interconnect the components of the equipment. The physical devices and wire list files together describe the physical characteristics of the hardware comprising the equipment.

## 1.2 Representation of Knowledge in the Knowledge Base 32

**[0048]** The knowledge contained in the knowledge base 32 is information from which the static view, dynamic view, and correlation therebetween, as well as the normality thereof, can be constructed. However, this knowledge must be understood by the inference engine

28. To allow the inference engine 28 to understand the information, it is represented as a "determinant", which represents a conceptual model of the equipment.

**[0049]** The determinant is abstractly represented as a graph consisting of nodes and edges. The nodes represent knowledge cells which contain knowledge useful for the inference engine 28 at a given point during diagnosis. Each knowledge cell may comprise, but is not limited to the following knowledge elements: (1) abstract operator label (i.e., the name of the knowledge cell); (2) cell reference (function definition, F); (3) physical device (substrate reference S(F)); (4) enablement condition of the substrate; (5) maintenance actions of the substrate; (6) downlink; (7) uplink - the downlink of other knowledge cells whose outputs come into this knowledge cell; (8) wire list associated with the substrate S(F); (9) normal operating state values; (10) logical state values; (11) special probe and inspection procedures; (12) leveler; (13) level jump condition; (14) test message; (15) CICH override; (16) GLCH override; and (17) cell type.

**[0050]** Each knowledge cell may comprise one of six different types. A blank cell is a knowledge cell that has no substrate but is used to group knowledge cells into categories to facilitate entry into the knowledge base 32 and to establish root cells described below. A regular cell is a knowledge cell that has a substrate S(F). A simple cell is a regular knowledge cell that cannot be decomposed, while a composite cell is a regular knowledge cell that may be decomposed. A logical cell is a regular knowledge cell that receives no outputs of any other regular cell, while a boundary cell is a regular knowledge cell which is a boundary of a taxonomical piece (i.e., it connects two or more taxonomical pieces). Each knowledge cell is constructed to have unique downlink, and two knowledge cells that have the same inputs are called associates.

**[0051]** There are four different connectors which are used as edges in forming the determinant. These are shown in FIG. 2. A link connector represents the circumstance where a regular cell sends its output to a regular cell. A leveler connector represents the condition where a composite cell is decomposable into a taxonomical piece consisting of other knowledge cells. An excursion connector is used to indicate that a given knowledge cell must be referred to in order to troubleshoot the substrate of a knowledge cell to which it originates. A blank edge is used to connect blank cells and therefore has no physical meaning. These four different types of connectors are utilized in FIGS. 3-5.

**[0052]** By representing knowledge elements as nodes and by using link, leveler, excursion, and blank edge connectors, the determinant of the equipment may be depicted graphically. A portion of one such graph is shown in FIG. 3. In this graph, knowledge cells 39 and 40 are connected by links to knowledge cell 42. The knowledge cell 42 in turn connected to knowledge cell 44 by a link, as well as to knowledge cell 46 by a leveler. Knowledge cell 48 is connected to the knowledge cell

46 by a link, while the knowledge cell 44 is connected to the knowledge cell 48 by an excursion. Knowledge cell 50 is connected to knowledge cell 46 by a link, while the knowledge cell 54 is connected to knowledge cell 50 by a blank edge. Finally, knowledge cell 54 is connected to the knowledge cell 52 by a link. By graphically representing the functional portions of the equipment in this manner, the functional interrelationship of its components may be characterized before being stored in the knowledge base 32.

**[0053]** An entire determinant representing the equipment may be represented as a lattice Klein surface as shown in FIG. 4. In FIG. 4, determinant 55 has knowledge cells 56-70. Knowledge cell 56 is connected to knowledge cell 58 by a leveler, while knowledge cell 58 is connected to knowledge cell 60 through a link. Knowledge cell 60 is connected to knowledge cell 62 through a link, while knowledge cell 64 is connected to knowledge cell 62 through a leveler. Knowledge cell 64 is connected to the knowledge cell 66 through a link, and knowledge cell 66 is in turn connected to knowledge cell 68 through a link. Knowledge cell 68 is connected to knowledge cell 70 through a link, while knowledge cell 70 is connected to knowledge cell 72 through a link. Finally, knowledge cell 72 is connected to knowledge cell 74 through a link, which is in turn connected to knowledge cell 56 through a link.

**[0054]** It will be apparent from FIG. 4 that knowledge cells 56 and 58 are on different levels since they are connected by a leveler. Accordingly, knowledge cell 56 may be decomposed into knowledge cells 58-62. In physical terms, this means that the component of the equipment represented by knowledge cell 56 is functionally dependent on three subcomponents represented by knowledge cells 58-62. By comparison, knowledge cell 56 is also connected to knowledge cell 64 through links via knowledge cells 66-74, which means that knowledge cell 56 is not functionally dependent on knowledge cell 64.

**[0055]** To minimize the time required for processing information regarding diagnosis of the equipment, the determinant 55 is normalized in the manner shown in FIGS. 5a-c. In FIG. 4a, knowledge cells A and B are on the same level on one path since they are connected by links through knowledge cell X. However, knowledge cell B is also at a lower level than knowledge cell A on a second path since knowledge cells A and B are also connected by a leveler. As shown in FIG. 5b, it is possible to cut the determinant 55 along knowledge cells X, Y and Z to form two taxonomical pieces TP 1 and TP 2, as shown in FIG. 5c. In the normalized graph shown in FIG. 4c, knowledge cell B is always at a lower level than knowledge cell A since it is no longer connected to knowledge cell A by a link through knowledge cell X.

**[0056]** Through this normalization process, the determinant 55 may be represented as a group of taxonomical pieces with all knowledge cells within a particular taxonomical piece connected by links. Each of the tax-

onomical pieces which comprise the determinant 55 are connected by levelers. The normalization process of the determinant 55 allows the determinant to be represented in such a way as to permit the knowledge system 24 to locate a fault in the equipment without examining each knowledge cell. In addition, by normalizing the determinant, the control over the inference engine 28 during diagnosis is somewhat easier as illustrated below.

### 1.3 Knowledge Engineering Requirements Regarding the Knowledge Base 32

**[0057]** To store information regarding the equipment in the knowledge base 32, there are three tasks which must be performed. First, the types of knowledge regarding the equipment which are needed for diagnosis must be determined. Secondly, the knowledge must be acquired from equipment documentation. Finally, the knowledge acquired from the documentation must be represented as a determinant in the manner described above. The results of these three steps may be transcribed to a set of data sheets, which data may then be entered into the knowledge base 32.

**[0058]** The specific sequence of tasks which must be performed to place the relevant knowledge in the knowledge base 32 is shown in FIG. 6. At steps 76 and 78, the user gathers design documentation regarding the equipment and obtains an understanding of the way the equipment operates. At step 80, the taxonomy of the knowledge cell is defined in that the functional components of the equipment are assigned to knowledge cells and the links between different knowledge cells are located. At step 82, an initial system model of the equipment is formed using the knowledge cells and links formed in step 80. An example of an initial system model is shown in FIG. 7. The half circles which are on the periphery of FIG. 7 are the boundary cells which are external to the system or subsystems on a consideration. The boundary cells A, B, and C whose outputs go into the intermediate system model are called in-boundary cells while boundary cells a, b, and c are called out-boundary cells in that their outputs go out of the intermediate system model.

**[0059]** Referring again to FIG. 6, in step 84, Giacomini process 1 is performed in which out-boundary cells are allowed to fuse into a single blank cell so as to become a root of a tree graph. As discussed above, a blank cell is an abstract cell which does not represent any specific function. By fusing the out-boundary cells of FIG. 7, FIG. 7 is transformed into the graph shown in FIG. 8. Knowledge cells B<sub>1</sub> and C<sub>1</sub>, which appear in FIG. 8 represent the associates of knowledge cells B and C respectively in that they have the same inputs.

**[0060]** As shown in FIG. 8, knowledge cell A appears twice on the same branch which indicates that a loop exists. Once all the loops in the transformed initial system model are detected, loops are graphically recreated on a tree graph as shown in FIG. 9. Giacomini process

2 is then performed on the discovered loops by introducing associate cells  $(1_0, 1_1)$ ,  $(2_0, 2_1, 2_2)$ ,  $(4_0, 4_1)$  and  $(5_0, 5_1)$  in the manner shown in FIG. 10. After transposing the graph along the line connecting the associate cells as shown in FIG. 10, the final nonhier-archical representation of the determinant is obtained as shown in FIG. 11.

**[0061]** Referring again to FIG. 6, at step 88, a hierarchical or normalized determinant is formed in the manner as shown in FIG. 12. Assume knowledge cell A represents a specific function and assume there exists a group of subfunctions within knowledge cell A which forms the determinant piece AD. If knowledge cell X represents the ROOT cell (i.e., the root of the tree graph representing the determinant) of the determinant piece AD, then a leveler exists between knowledge cell A and knowledge cell X. Examining other knowledge cells in the same manner, each function of the equipment is decomposed into subfunctions to form a set of determinant pieces which are connected by levelers.

**[0062]** At steps 90 and 92, knowledge cells are connected and normal state values are defined and acquired. In this step, each knowledge cell that appears in the determinant model is assigned the types of knowledge described above. At steps 94 and 96, the determinant, together with the information collected in steps 90 and 92, are transcribed into formatted data sheets and inputted through learning module 36 into the knowledge base 32. At steps 98 and 100, the "goodness" of the knowledge base 32 is determined hierarchically by using three criteria. The first criterion is the analytical truth of the information contained in the knowledge base 32. To determine the analytical truth of the information, the knowledge base 32 is examined by a knowledge consistency checker within the learning module 36 to determine whether there are any logical contradictions.

**[0063]** The knowledge base 32 is then examined to determine whether the analytical truth corresponds to reality. This is done by determining the synthetic truth of the knowledge base 32. The synthetic truth is established to the extent that a set of selected faults inserted into the equipment are correctly isolated. For those selected faults which are not correctly fault isolated, a revision of the knowledge base 32 is performed.

**[0064]** The third goodness criteria against which the knowledge base 32 is tested is completeness. Completeness means that the knowledge base 32 has a complete knowledge about the manner in which the electronic equipment 32 is designed to work. The degree of completeness is determined by the frequency of ignorance or inability of the system 24 to isolate a fault. After the completeness of the knowledge base 32 has reached a level acceptable to the user (for example, 99 fault isolations out of 100 attempts), the knowledge base 32 may be used in conjunction with the inference engine 28 to perform diagnosis.

## 2. INFERENCE ENGINE 28

**[0065]** Before discussing the specific operation of the inference engine 28, several terms shall be defined. The first is the term "cyber". Since the determinant of the equipment is represented as a collection of taxonomical pieces, a given knowledge cell A must be in some taxonomical piece X. Accordingly, the location of a knowledge cell may be represented as the ordered pair  $(X,A)$ , where A represents the name of the knowledge cell and X represents the name of the taxonomical piece containing the knowledge cell A. The ordered pair  $(X,A)$  is referred to as a cyber. Accordingly, if f is a fault that is caused by an equipment malfunction, f must lie in some cyber  $(X,A)$ . A cyber  $(X,A)$  which contains a fault f is called a "fault cell", and the knowledge cell associated with the cyber is called a faulty knowledge cell.

**[0066]** As more fully discussed below, the inference engine 28 examines a sequence of cybers in locating a fault in the equipment. The specific sequence of cybers which the inference engine 28 examines is called the cyber transition sequence. In general, the inference engine 28 generates a cyber transition sequence according to one of the following rules: (1) the inference engine 28 chooses a cyber  $(X,B)$  such that the taxonomical piece X is bisected so that approximately one half of the cybers which do not contain fault cells in the taxonomical piece can be deleted from further consideration; (2) if a cyber  $(Y,B)$  is a boundary of the taxonomical piece Y, then the taxonomical piece Y which is connected to the taxonomical piece X is entered and then the inference engine 28 moves to cyber  $(Y,B)$ ; (3) if the knowledge cell A of the cyber  $(X,A)$  is a composite knowledge cell, then the inference engine 28 "levels down" to a taxonomical piece Y and then chooses a cyber  $(Y,B)$  such that the taxonomical piece Y can be bisected, or (4) the inference engine 28 may directly jump to any cyber  $(Y,B)$  if told to do so by means of the GLCH Override described below.

**[0067]** If the inference engine 28 finds a fault cell indicating a faulty knowledge cell A while examining the cybers in the cyber transition sequence, then the inference engine 28 locates the fault in that particular knowledge cell A. Accordingly, the control of the inference engine 28 either ends at a fault cell if the fault cell contains a fault, or it constructs another cyber transition sequence. The sequence in moving from one fault cell to a cyber transition sequence of another fault cell is called a fault-transition sequence.

**[0068]** Before the flow charts associated with the inference engine 28 are described, a general description of the behavior of the inference engine 28 during fault diagnosis will be presented. Assume that the equipment malfunctions and the cause of the malfunction is the fault F. Then the normality criteria of the cyber in the determinant must be violated. If the fault cell of the taxonomical piece  $(X,A)$  is found, the fault F must reside either in the underlying physical device (i.e., the sub-



strate S(A)), in its enablement conditions, in its maintenance discipline (i.e., maintenance actions), or in the wires associated with the substrate S(A) (i.e., open or pinched).

**[0069]** Accordingly, the inference engine 28 can adopt either one of four possible behaviors. The inference engine 28 can find a fault cell indicating that the cyber transition sequence is finite. This condition is represented by a GLCH-Success signal. Alternatively, the inference engine 28 may not be able to find a fault cell indicating that the cyber transition sequence forms a loop and is therefore infinite. When the inference engine 28 cannot find a fault cell, the inference engine 28 terminates diagnosis with the statement to the user that the inference engine 28 is executing an infinite loop. In doing so, the system 24 lists all the cybers in the loop as well as prints out an error in the system taxonomy and the taxonomical pieces involved. The occurrence of this condition is indicated by a GLCH-Failure signal.

**[0070]** Further, the inference engine 28 may find a fault in the fault cell, indicated by a CICH-Success signal, or it may not find a fault in all fault cells in the taxonomy indicated by a CICH-Failure signal. When the inference engine 28 cannot find a failure in all fault cells, it terminates the diagnosis and informs the user that the knowledge cell associated with the cyber (X,A) has insufficient knowledge and that additional knowledge is needed. Accordingly, the inference engine 28 is able to determine what it does not know.

**[0071]** The inference engine 28 contains six logical components. The first logical component is the global control heuristic ("GLCH"). The global control heuristic generates the cyber transition sequences, and finds a fault cell or terminates diagnosis indicating that the system 24 is ignorant as to the specific type of failure. The second logical component is the blackboarder-taxonomy routine ("BBDR-TAX"). Because the global control heuristic acts only on the determinant and not on the knowledge base 32 directly, the blackboarder-taxonomy routine accesses the knowledge base 32 and constructs a taxonomical piece of the determinant upon which the global control heuristic acts. The third logical component is the blackboarder-knowledge cell routine ("BBDR-KC") which is used to retrieve the various knowledge types which are associated with a particular knowledge cell. These include uplinks, downlinks, the normal operating state values and so forth.

**[0072]** The fourth logical component is the cell internal control heuristic ("CICH") which acts on the knowledge cell that the blackboarder-knowledge cell routine has constructed. The cell internal control heuristic isolates a fault in the knowledge cell or decomposes the knowledge cell if the cell is a composite. The cell internal control heuristic then calls the global control heuristic to restart a separate cyber transition sequence or cause the diagnosis to exit with a cell knowledge deficiency statement.

**[0073]** The fifth logical component, the leveler-in-

struction routine (LEV-INS), tells the blackboarder-taxonomy cell routine what portion of the knowledge base 32 to retrieve and causes the construction of a determinant piece upon which the global control heuristic and the cell internal control heuristic act. Although the levelerinstruction routine, blackboarder-taxonomy and blackboarder-knowledge cell routines form part of the blackboarder 30, they will be discussed in conjunction with the inference engine 28 for purposes of simplicity.

**[0074]** Finally, the goodness propagation routine ("GP") is used to allow quick location of fault cells. As the global control heuristic generates cybers in the cyber transition sequence, the domain of the search in the taxonomical piece is cut approximately in half. In order to locate fault cells in the determinant, it is necessary to know which half of the determinant contains the fault cell and which half does not. Accordingly, the goodness propagation routine extracts a maximum of information from a "single probe" to propagate "goodness" to a maximum number of yet untested knowledge cells which the inference engine 28 will not have to examine during diagnosis. The goodness propagation routine will be discussed below in conjunction with the global control heuristic.

**[0075]** A block diagram of the inference engine 28 is shown in FIG. 13. The inference engine 28 comprises a find fault cell routine ("FF-C") 102 and a find fault routine ("FF") 104. The find fault cell routine 102 is used to generate the cyber transition sequence which terminates with a cyber having the name of a fault cell. The find fault routine 104 is then used to isolate the fault within the fault cell. During operation, if the find fault cell routine 102 locates a fault cell, the processing flows along path "a" to the find fault routine 104 with the value of task i equal to zero. When the task i equals zero, the find fault routine 104 is instructed to check for each possible fault starting from the initial task (i.e., the find fault routine 104 will examine the downlinks, uplinks, enablement conditions and so forth to determine whether the equipment is exhibiting these particular faults). If it is determined that the knowledge cell under consideration is in fact a composite knowledge cell, the find fault cell routine 102 is instructed to level down to another taxonomical piece. This is done by the generation of a TT=0 output by the find fault routine 104 which is delivered to the find fault cell routine 102 along path "d". If it is determined that the fault of the current knowledge cell is due to another cyber, the find fault cell routine 102 is instructed to generate another cyber transition sequence beginning from another knowledge cell. This is done by generation of a TT=0 output by the find fault routine 104 which is delivered to the find fault cell routine 102 along path "e".

**[0076]** If a fault cell is not located, the find fault routine 104 is instructed to search task i other than task i = 0 which is done through path "b". If the find fault cell routine 102 is instructed to search for a fault cell in a particular taxonomical piece and that taxonomical piece

does not contain a fault cell, this information is returned to the find fault routine 104 through path "c" which sets task i unequal to O. This allows the find fault routine 104 to terminate the diagnostic session with an "F-cell ignorance" message indicating that the system 24 does not contain information regarding the fault cell. This is shown in FIG. 13 at step 112. Further, if it is determined that all knowledge cells have been searched without locating the fault cell, the diagnostic session terminates with a "loop ignorance" message as shown in step 114. If the find fault routine 104 is able to locate the fault within a particular fault cell, this fact is displayed to the user at step 116.

**[0077]** A more complete description of the find fault cell routine 102 will now be made with reference to FIG. 14. The find fault cell routine 102 starts at step 122 which calls the leveler-instruction routine described below to obtain the identification of the taxonomical piece if the next cyber in the cyber transition sequence is located. In doing so, the routine passes information to the leveler-instruction routine at step 122 regarding the identification of the present determinant piece D, the present knowledge cell A, whether the find fault cell routine 102 was accessed by the find fault routine 104 in either an override condition (TT=O), through a boundary cell (TT=B), or by levelling down to another taxonomical piece (TT=0). Finally, the routine also passes information to the leveler-instruction routine at step 122 regarding the task i which is being presently examined.

**[0078]** At step 124 it is determined whether or not the taxonomical piece under consideration contains a fault cell. This is done by examining the contents of the level jump conditions file. If the taxonomical piece under consideration does not contain a fault cell, control is passed to the find fault routine 104 at step 126.

**[0079]** If at step 124, it is determined that the taxonomical piece under consideration contains a fault cell, step 128 is executed, and the blackboarder-taxonomy routine is accessed. As described below, the blackboarder-taxonomy routine accesses the taxonomy file, the boundary file, and the test message file of the knowledge base 32 and uses this information to construct a determinant piece on the blackboard 34. In executing step 128, the find fault cell routine 102 passes information to the blackboarder-taxonomy cell routine regarding the label of the taxonomical piece under consideration, as well as the test message number associated with the taxonomical piece.

**[0080]** After executing step 128, step 130 is executed which calls the global control heuristic which in turn searches for a fault cell until it finds the taxonomical piece containing the fault cell or passes through boundary cells into a new taxonomical piece. In executing step 130, information is passed to the global control heuristic regarding the label of the current taxonomical piece, and the test message number of the taxonomical piece.

**[0081]** At step 132, it is determined whether it is necessary to move to a new taxonomical piece to locate the

fault cell. If a new taxonomical piece must be examined to find the fault cell, step 122 is executed as described above. If at step 132, it is determined that it is not necessary to move to another taxonomical piece, step 134 is executed which determines whether the fault cell has been located. If the fault cell has not been located, the inference engine 28 then executes the find fault routine 104 by way of the step 136. If at step 134 it is determined that the fault cell has been found, the inference engine 28 then executes the find fault routine 104 with task i equal to 0 by way of step 138.

**[0082]** The find fault routine 104 will now be described with reference to FIG. 15. At step 140, the value of task i is generated by find fault cell routine 102 is read. If at step 142, it is determined that task i equals 0, indicating that a fault cell has been found, the blackboarder-knowledge cell routine is called at step 144 which constructs the complete knowledge cell on the blackboard 34. In doing so, the find fault routine 104 passes to the blackboarder-knowledge cell routine information regarding the name of the fault cell and taxonomical piece in which the fault cell is located.

**[0083]** After executing step 144, or if at step 142 task i does not equal zero, step 146 is executed which calls the cell internal control heuristic. As described below, the cell internal control heuristic is used for isolating a fault in the fault cell. In executing step 146, the find fault routine 104 passes to the cell internal heuristic the name of the taxonomical piece in which the knowledge cell is located, the name of the fault cell, and the value of task i.

**[0084]** If the cell internal control heuristic succeeds in isolating the fault within the fault cell, there is no need to construct another cyber transition sequence. Accordingly, the inference engine 28 then calls the find fault cell routine 102 at steps 148 and 150 of the find fault routine 104. In this case, the find fault cell routine 102 will either decompose the current fault cell if it is a composite cell by leveling down (i.e., TT=0), or will terminate diagnosis with an "F-cell ignorance" signal at step 112. If the cell internal control heuristic does not fault isolate, and indicates that a new cyber transition sequence must be generated, the find fault cell routine 102 is called at step 152. When this occurs, the value of TT is set equal to zero to indicate that an excursion departure will occur in indicating that the fault of the knowledge is due to a fault in another cyber.

**[0085]** The blackboarder-taxonomy cell routine will now be described with reference to FIG. 16. The blackboarder-taxonomy routine, generally designated by the numeral 154, has an initial step 156 in which both the name of the taxonomical piece under consideration and the test message number associated with the taxonomical piece are read. At step 158, it is determined whether the particular taxonomical piece under consideration is on the blackboard 34. If it is not on the blackboard 34, step 160 is executed which retrieves the taxonomical piece from the taxonomy the boundary files.

**[0086]** After executing step 160, or if at step 158 it is

determined that the taxonomical piece is on the blackboard 34, step 162 is executed. At step 162, a test message is retrieved from the test message file and the test message is displayed on the user interface 26. After executing step 162, step 164 is executed which copies the taxonomical piece onto a portion of the blackboard 34. Control is then returned to the calling step at return step 166.

**[0087]** The leveler-instruction routine will now be described with reference to FIG. 17. The leveler-instruction routine, generally designated by the numeral 168, comprises an initial step 170 which reads the name of the current determinant piece D, the value of TT, and the value of task i. At step 172, the determinant piece is retrieved from the determinant directory file together with the name of the determinant piece, the level jump condition number, and the test message number. At step 174, the level jump condition is retrieved from the level jump conditions file. At step 176, a level jump condition is performed in which the user is instructed to perform a test at user interface step 178. If the level jump condition is passed as determined at step 180, step 182 is executed which instructs the inference engine 28 to execute step 128 of the find fault cell routine 102. In doing so, step 182 passes an indication that the level jump condition was passed, and passes information regarding the test message number and the label of the taxonomical piece under consideration to the find fault cell routine 102.

**[0088]** If at step 180, it is determined that the level jump condition was not passed, it is determined whether the find fault routine 104 instructed the find fault cell routine 102 to level down (i.e., whether the value of TT=0). If at step 184 it is determined that the find fault routine 104 requested the find fault cell routine 102 to level down, step 186 is executed which instructs the inference engine 28 to execute step 126 of the find fault cell routine 102.

**[0089]** If at step 184 it is determined that the find fault routine 104 did not instruct the find fault cell routine 102 to level down, step 188 is executed. At step 188, a warning is issued to the user that the find fault cell routine 102 was asked to determine whether a particular determinant piece contains a fault, but the find fault routine 104 has already determined that the particular determinant piece did not contain a fault.

**[0090]** Step 190 is then executed in which the user is asked at step 192 whether the diagnosis should continue. If the user instructs the leveler-instruction routine 168 to continue, step 194 is then executed which causes the inference engine 28 to execute step 128 of the find fault cell routine 102. In doing so, the step 194 indicates to step 128 that the level jump condition was passed by override, and passes information regarding the taxonomical piece label and the test message number. If the user decides to terminate the diagnostic session after a warning has issued at step 188, step 196 is executed which instructs the inference engine 28 to execute step

126 of the find fault cell routine 102.

**[0091]** The blackboarder-knowledge cell routine will now be described with reference to FIG. 18. The blackboarder-knowledge cell routine, generally designated by the numeral 198, starts at step 200 wherein the identification of the current taxonomical piece and current knowledge cell in the taxonomical piece is read. In addition, the test message number associated with the taxonomical piece and the knowledge cell is also read. At step 202, uplinks associated with the knowledge cell are retrieved from the taxonomical piece. At step 204, the downlinks from the downlink file associated with the knowledge cell. After executing step 204, step 206 is executed in which the wire lists or logical downlink states associated with the knowledge cell are retrieved. At step 208 the normal operating state values associated with the knowledge cell are retrieved from the normal operating state values file.

**[0092]** Step 210 is then executed which retrieves information regarding special probe/inspection procedure data of the knowledge cell from the special probe/inspection procedure file. At step 212, enablement conditions associated with the knowledge cell are retrieved from the enablement conditions file. Step 214 is then executed which retrieves all maintenance actions regarding the knowledge cell from the maintenance actions file. At step 216, override conditions of the global control heuristic and the cell internal control heuristic are retrieved. At step 218, levelers associated with the knowledge cell are retrieved from the leveler file. Step 220 is then executed which retrieves physical device information associated with the knowledge cell from the physical device file. The inference engine 28 then executes step 128 of the find fault cell routine 102 by way of return step 222.

**[0093]** The global control heuristic will be described with reference to FIGS. 19a-c. The global control heuristic, generally designated by the numeral 224, starts at step 226 which reads the name of the particular knowledge cell in question, the associated taxonomical piece and the value of TT and downlink wires. Step 228 is then executed in which a variable ROOT is set equal to the given taxonomical piece. At step 230, the present value of TT is determined. If the value of TT equals 0 indicating that the find fault cell routine 102 was executed due to an override condition, step 232 described below is executed. If at step 230, it is determined that the value of TT is equal to B thereby indicating that the find fault cell routine 102 was accessed through a boundary cell, step 234 described below is executed.

**[0094]** If at step 230, it is determined that the value of TT=0 thereby indicating that the find fault cell routine 102 was accessed through a leveling down condition, step 236 is executed. At step 236, all knowledge cells in the taxonomical piece are assigned integer values from 0 to N, where zero is always assigned to the ROOT cell. After executing step 236, step 238 determines whether the value of N equals zero. If the value of N

equals zero, at step 240, it is determined whether ROOT has changed during processing. If the ROOT has not changed during processing, this indicates that the determinant piece does not contain a knowledge cell with a fault, and step 242 causes the inference engine 28 to execute step 136 of the find fault cell routine 102.

**[0095]** If at step 240, it is determined that ROOT has changed during processing, step 244 is executed which determines whether ROOT represents a boundary cell. If ROOT represents a boundary cell, the inference engine 28 is instructed to review a different knowledge cell in a different taxonomical piece. In doing so, the inference engine 28 sets the value of TT=B, as well as identifies wire i which is the downlink of the boundary cell whose output is abnormal. The global control heuristic 224 then instructs the inference engine 28 to execute step 122 of the find fault cell routine 102 at step 246.

**[0096]** If at step 244, it is determined that ROOT does not represent a boundary cell, step 248 sets task i equal to zero which indicates that the fault cell has been found. The global control heuristic 224 then instructs the inference engine 28 to execute step 138 of the find fault cell routine 102 at return step 250.

**[0097]** If at step 238, it is determined that the value of N does not equal zero, step 252 is executed, wherein integers assigned in step 238 are binary-dissected, rounding upward if necessary. For purposes of this discussion, it will be assumed that the knowledge cell whose assigned integer number is in the middle of the binary dissection is knowledge cell B which is located in taxonomical piece X. After executing step 252, step 254 assigns a value to the knowledge cell under consideration equal to the value of knowledge cell B derived from the binary dissection. After executing step 254, step 256 probes the downlink associated with knowledge cell B.

**[0098]** Step 258 is then executed which determines whether the outputs associated with the knowledge cell B are normal. If it is determined that the outputs associated with knowledge cell B are not at their appropriate levels, step 260 is executed which determines whether the downlink which was abnormal was a logical type or wire type downlink. If it is determined at step 260 that the downlink was a wire type downlink, step 262 retains information concerning the wire that had a faulty output. After executing step 262, or if at step 260 it is determined that the downlink was a logical type, step 264 is executed in which the value of ROOT is set equal to B. After executing step 264, step 266 eliminates those knowledge cells which do not have faulty outputs. After executing step 266, a decision is made at step 236.

**[0099]** If at step 268, the downlinks do not indicate that the current knowledge cell is faulty, step 268 is executed, which indicates if the current knowledge cell, as well as all knowledge cells which provide outputs to the current knowledge cell are not faulty. This process is called goodness propagation. After goodness propagation occurs at step 268, step 270 is executed, in which knowledge cells which have been determined not to be faulty

are eliminated from consideration. After executing step 270, step 272 determines whether there is a global control heuristic override associated with the current knowledge cell. If there is no global control heuristic override step 238 is executed.

**[0100]** If at step 272, it is determined that there exists a global control heuristic override associated with the current knowledge cell, step 274 is executed which determines whether the determinant piece to which the global control heuristic override is directed is the current determinant piece. If it is not, step 276 changes the determinant piece to that which was the subject of the global control heuristic override by instructing the inference engine 28 to execute step 126 of the find fault cell routine 102.

**[0101]** If it is determined that the current determinant piece is the same as the determinant piece to which the global control heuristic override is directed, step 278 determines whether goodness propagation has been performed on the current knowledge cell. If it is determined that goodness propagation has already occurred with respect to the current knowledge cell, a warning is issued at step 280 to inform the user that the current knowledge cell was already found not to be a fault cell. After executing step 280, the user is asked whether the diagnosis should continue at step 282. If the diagnosis is not to continue, the inference engine 28 is instructed to execute step 128 of the find fault cell routine 102. If the user determines at step 282 that the diagnosis should continue, or at step 278 it is determined that goodness propagation has not occurred with respect to the current knowledge cell, step 284 is executed.

**[0102]** If at step 230, the value of TT is found to be equal to B, which indicates that the find fault cell routine 102 was accessed through a boundary cell, the user is asked to probe wire i at step 286. If it is determined that the output of wire i is abnormal, step 264 is executed at step 286. If the output from wire i is normal, step 288 is executed which issues a warning to the user that, while the present output wire i is normal, its past output was abnormal. If the user determines that diagnosis should continue, step 236 is executed by means of step 290. If the user decides diagnosis should be discontinued, diagnostic activity is terminated at step 292 and knowledge base error correction is requested.

**[0103]** If at step 230, it is determined that the value of TT equals O thereby indicating a global control heuristic override has occurred, step 294 is executed which probes the downlinks associated with the knowledge cell. If one of the downlinks is abnormal, step 264 is executed by means of step 296. If the downlinks associated with the knowledge cell are normal, step 298 is executed which issues a warning indicating that the global control heuristic 224 was instructed to examine a knowledge cell for a fault but the knowledge cell has been determined not to be faulty. If the user decides to continue with the diagnosis, step 236 is executed by means of step 300. If the user decides to terminate diagnosis, di-

agnostic activity is terminated at step 302 and knowledge base error correction is requested.

**[0104]** The cell internal control heuristic will now be described with respect to FIGS. 20a-c. The cell internal control heuristic, generally designated by the numeral 304, starts at step 306 in which the current knowledge cell, together with its taxonomical piece, the value of TT, and the wire i are read. Step 308 is then executed which determines whether the value of task i is equal to zero. If the value of task i is equal to zero, step 310 is executed which determines whether the downlink of wire i is equal to zero, indicating that the downlink is a wire type. If at step 310 it is determined that the downlink of wire i is not equal to zero, step 312 is executed which instructs the user to detach the wire connected to the substrate represented by the knowledge cell A.

**[0105]** After executing step 312, the user is then asked at step 312 whether output is normal after the wire i has been detached. If the user indicates that a normal output has appeared, step 316 provides a message that indicates that the wire i has been pinched and that isolation has occurred. If the output of wire i is not normal, or if at step 310 the downlink of wire i is equal to 0, it is determined at step 318 whether there exists a cell internal control heuristic override associated with the knowledge cell. If such an override exists, step 320 is executed. If no cell internal control heuristic override exists, step 322 is executed in which the sequence of tasks to be exhausted are selected.

**[0106]** After executing step 322, or after modifying the intrinsic task sequence at step 230, or if at step 308 it is determined that the task i does not equal 0, step 324 is executed, which brings in the next task as task i. After executing step 324, it is determined at step 326 whether there are any more tasks to consider. If there are no more tasks to consider, step 328 determines whether the knowledge cell is simple. If it is, step 330 is executed which indicates that the fault appears in the particular physical device and that fault isolation has occurred. If the cell is not simple, step 332 is executed which terminates diagnosis with an ignorance condition indicating that the composite cell has not been found to contain a fault.

**[0107]** If at step 326, it is determined that there are more tasks to consider in the diagnosis, step 334 determines whether task i is an uplink check. If it is determined that task i is an uplink check, step 336 asks the user to measure the output of the wire leading to the substrate represented by the knowledge cell A in order of talley (i.e., which is determined by the inference engine 28 as having the highest probability of failure).

**[0108]** Next, it is determined whether the outputs from the uplink wires of the knowledge cell under examination are normal at step 338. If all the outputs are normal, step 324 is executed. If at step 338 the outputs from the uplink wires associated with the current knowledge cell are abnormal, step 340 is executed in which an open wire fault is indicated.

**[0109]** If at step 334 the task i does not represent an uplink check, step 342 determines whether the task i represents an enablement conditions check. If task i represents an enablement conditions check, the enablement conditions are probed at step 344. After the enablement conditions are probed, it is determined at step 346 whether all the enablement conditions are normal. If all the enablement conditions are normal, step 324 is executed. If at step 346 it is determined that not all the enablement conditions probed are normal, step 348 is executed which indicates that an enablement conditions fault has occurred and that the fault has been isolated.

**[0110]** If at step 342, it is determined that task i does not represent an enablement conditions check, step 350 is executed, which determines whether task i represents a maintenance actions check. If step 352 requests the user to take the appropriate maintenance action. If the user indicates that step 354 that all maintenance actions have been taken, step 324 is executed. If all maintenance action have not been taken, step 356 determines whether there has been a permanent excursion departure (i.e., that the inference engine 28 should examine a different cyber). If there has been no permanent excursion departure, a maintenance actions fault is indicated and diagnosis is terminated at step 358. If at step 356 it is determined that there has been a permanent excursion departure, step 360 is executed in which the inference engine 28 causes execution of step 152 by way of step 360.

**[0111]** If at step 350 it is found that task i does not represent a maintenance action check, step 362 determines whether task i represents a level down condition. If it does, step 36 causes the inference engine 28 to execute step 150. If task i does not represent a level down condition, step 366 instructs the find fault cell routine 102 to level down.

**[0112]** It should be understood that while the present invention was described in connection with one specific embodiment, other modifications will become apparent to one skilled in the art upon a study of the specification, drawings and following claims.

## Claims

1. An expert system (24) for diagnosing faults in equipment, said system including a computer having a memory for storing diagnostic routines, a knowledge base (32) accessible by the computer, an inference engine (28) which interfaces with the knowledge base, and a user interface (26) which permits communication between the expert system and an operator,

said knowledge base (32) characterizing the physical, functional and operational relationships of the components and between components of the equipment and information con-

cerning selected operating conditions of the equipment such that said knowledge base (32) contains a conceptual model of the equipment; and

said inference engine (28) interfacing with the operator by way of the user interface and comprising recursive diagnostic routines which operate on the data contained in the knowledge base (32) to generate symptom diagnostic rules in response to operator entered symptoms which are indicative of the faults present in the equipment such that each symptom diagnostic rule correlates a symptom to its fault and is indicative of the fault;

characterized in that

said knowledge base (32) represents the system taxonomy or determinant (Figure 4: 55) of the equipment, said determinant having knowledge cells (56 - 70), a knowledge cell representing a defined function (F) which a component (substrate S(F)) of the equipment performs, a knowledge cell may be connected to another knowledge cell through a downlink (58 - 60) if both cells are on the same level, or through a leveler (56 - 58) if both cells are on different levels, said determinant (55) being partitioned into taxonomical pieces, each taxonomical piece relating to a subset of knowledge cells connected by links, said knowledge base comprising the following files:

(1) a taxonomy file containing the name of each knowledge cell, the cell's downlinks, and the number of levelers;

(2) a leveler file containing leveler information indicating whether knowledge cells are composite cells which may be decomposed into smaller knowledge cells, so indicating whether a component of the equipment may be decomposed into a plurality of smaller components;

(3) a boundary file containing a listing of knowledge cells which are also boundary cells in each taxonomical piece;

(4) a wire list file containing information regarding the wires which interconnect the components of the equipment;

(5) a level jump conditions file containing information for each taxonomical piece regarding a test procedure to determine whether a faulty cell exists within the taxonomical piece;

(6) a test message file for storing outputs of all functions in a taxonomical piece which are generated when a test message representing a stimulus is applied to the components represented by said taxonomical piece;

(7) an enablement condition file containing information regarding conditions that must be satisfied by the components of the equipment for proper operation;

(8) a normal operational state values file containing information regarding correct outputs of a component when the stimulus is applied thereto;

wherein for locating a faulty component the inference engine examines a specific sequence (cyber transition sequence) of knowledge cells or cybers (a cyber (X,A) is a knowledge cell A in a taxonomical piece X), wherein the inference engine moves to a further cyber in dependence whether a cell examined by the inference engine is a boundary cell or a composite cell.

2. The expert system according to claim 1 further characterized in that said knowledge base (32) further comprises files containing maintenance actions taking into account a complex set of factors which are encountered through experience, which actions may be applied to said equipment.
3. The expert system according to claim 2 further characterized in that each of said knowledge cells is uniquely identified by a cyber representing the location of a knowledge cell within the determinant, and said inference engine (28) includes means for generating a cyber transition sequence.
4. The expert system according to claim 2 further characterized in that said expert system (24) further comprises a blackboard (34) coupled between said knowledge base (32) and said inference engine (28) and means for accessing the data stored in the knowledge base (32) and retrieving data for constructing a portion of said determinant on said blackboard for the inference engine (28) to perform its diagnosis.
5. The expert system according to claim 4 further characterized in that said means for constructing a portion of said determinant on said blackboard (34) comprises a blackboarder (30), said inference engine (28) comprises means for indicating to said blackboarder which portion of said conceptual model to retrieve.

## Patentansprüche

1. Expertensystem (24) zum Diagnostizieren von Fehlern in einem Gerät, wobei das System aufweist einen Computer mit einem Speicher zum Speichern von Diagnoseroutinen, eine Wissensbasis (32), auf die der Computer zugreifen kann, eine Inferenzmaschine (28), die über eine Schnittstelle verbunden ist mit der Wissensbasis, und eine Benutzerschnittstelle (26), die eine Kommunikation zwischen dem Expertensystem und einer Bedienperson gestattet,

wobei die Wissensbasis (32) die physikalischen, funktionalen und betrieblichen Beziehungen der Komponenten und zwischen Komponenten des Gerätes und von Information betreffend ausgewählte Betriebszustände des Gerätes charakterisiert, derart, daß die Wissensbasis (32) ein konzeptionelles Modell des Gerätes enthält; und

wobei die Inferenzmaschine (28) mittels der Benutzerschnittstelle mit der Bedienperson in Verbindung steht und rekursive Diagnoseroutinen enthält, die mit den Daten arbeiten, die in der Wissensbasis (32) enthalten sind, um in Antwort auf die Eingabe von Symptomen durch die Bedienperson Symptomdiagnoseregeln zu erzeugen, wobei die Symptome in dem Gerät vorliegende Fehler anzeigen, derart, daß jede Symptomdiagnoseregeln ein Symptom mit seinem Fehler korrelativ in Verbindung setzt und den Fehler anzeigt;

dadurch gekennzeichnet, daß die Wissensbasis (32) die Systemtaxonomie oder -determinante (Fig. 4: 55) des Gerätes darstellt, wobei die Determinante Wissenszellen (56-70) aufweist, wobei eine Wissenszelle eine definierte Funktion (F) repräsentiert, die eine Komponente (Substrat S(F)) des Gerätes durchführt, wobei eine Wissenszelle mit einer anderen Wissenszelle über einen Abwärtslink (58-60), wenn sich beide Zellen auf derselben Ebene befinden, oder mittels eines Levelers (56-58) verbunden werden kann, wenn sich die beiden Zellen auf unterschiedlichen Ebenen befinden, wobei die Determinante (55) in taxonomische Teile unterteilt ist, wobei jedes taxonomische Teil sich auf eine Untermenge von Wissenszellen bezieht, die durch Links miteinander verbunden sind, wobei die Wissensbasis die folgenden Dateien aufweist:

(1) eine Taxonomiedatei, die den Namen von jeder Wissenszelle, die Abwärtslinks der Zelle und die Anzahl von Levelern enthält;

(2) eine Levelerdatei, die Levelerinformation enthält, die anzeigt, ob Wissenszellen Kompo-

sitzellen sind, die in kleinere Wissenszellen zerlegt werden können, und somit anzeigt, ob eine Komponente des Gerätes in eine Vielzahl von kleineren Komponenten zerlegt werden kann;

(3) eine Grenzdatei, die eine Liste von Wissenszellen enthält, die in dem jeweiligen taxonomischen Teil auch Grenzzellen sind;

(4) eine Verdrahtungslistendatei, die Information enthält betreffend die Verdrahtungen, die die Komponenten des Gerätes untereinander verbinden;

(5) eine Ebenensprungbedingungsdatei, die für jedes taxonomische Teil Information betreffend eine Testprozedur enthält, um zu bestimmen, ob innerhalb des taxonomischen Teils eine fehlerhafte Zelle existiert;

(6) eine Testmitteilungsdatei zum Speichern von Ausgaben aller Funktionen in einem taxonomischen Teil, die erzeugt werden, wenn eine Testmitteilung, die einen Stimulus darstellt, an die Komponenten angelegt wird, die durch das taxonomische Teil repräsentiert sind;

(7) eine Freigabebedingungsdatei, die Information enthält betreffend Bedingungen, die für einen ordnungsgemäßen Betrieb von den Komponenten des Gerätes erfüllt werden müssen;

(8) eine Datei mit Zustandswerten für Normalbetrieb, die Information betreffend korrekte Ausgänge einer Komponente enthält, wenn ein Stimulus daran angelegt wird;

wobei die Inferenzmaschine zum Lokalisieren einer fehlerhaften Komponente eine spezifische Sequenz (Cyberübergangssequenz) von Wissenszellen oder Cybern (ein Cyber (X,A) ist eine Wissenszelle (A) in einem taxonomischen Teil X) prüft, wobei die Inferenzmaschine zu einem weiteren Cyber übergeht in Abhängigkeit davon, ob eine von der Inferenzmaschine geprüfte Zelle eine Grenzzelle oder eine Kompositzelle ist.

2. Expertensystem nach Anspruch 1, ferner dadurch gekennzeichnet, daß die Wissensbasis (32) Dateien aufweist, die Wartungsaktionen enthalten, die einen komplexen Satz von Faktoren berücksichtigen, die man durch Erfahrung erfährt, wobei diese Aktionen an dem Gerät ausgeübt werden können.

3. Expertensystem nach Anspruch 2, ferner dadurch gekennzeichnet, daß jede der Wissenszellen eindeutig durch einen Cyber identifiziert ist, der den Ort

einer Wissenszelle innerhalb der Determinante repräsentiert, und wobei die Inferenzmaschine (28) Mittel zum Erzeugen einer Cyberübergangssequenz enthält.

4. Expertensystem nach Anspruch 2, ferner dadurch gekennzeichnet, daß das Expertensystem (24) ferner ein Blackboard (34) aufweist, das zwischen der Wissensbasis (32) und der Inferenzmaschine (28) angeschlossen ist, und Mittel aufweist zum Zugreifen auf die in der Wissensbasis (32) gespeicherten Daten und zum Holen von Daten zum Aufbauen eines Abschnittes der Determinante auf dem Blackboard für die Inferenzmaschine (28), damit diese ihre Diagnose durchführen kann.
5. Expertensystem nach Anspruch 4, ferner dadurch gekennzeichnet, daß die Mittel zum Aufbauen eines Abschnittes der Determinante auf dem Blackboard (34) einen Blackboarder (30) aufweisen, wobei die Inferenzmaschine (28) Mittel aufweist, um dem Blackboarder anzuzeigen, welcher Abschnitt des konzeptionellen Modells zu holen ist.

#### Revendications

1. Système expert (24) de diagnostic d'anomalies dans un équipement, ledit système comportant un calculateur ayant une mémoire pour stocker des programmes de diagnostic, une base de connaissances (32) accessible par l'ordinateur, un moteur d'inférence (28) qui forme une interface avec la base de connaissances, et une interface utilisateur (26) qui permet une communication entre le système expert et un opérateur,

ladite base de connaissances (32) caractérisant les relations physiques, fonctionnelles et opérationnelles des composants et entre les composants de l'équipement et les informations concernant les conditions de fonctionnement sélectionnées de l'équipement, de sorte que cette base de connaissances (32) contient un modèle conceptuel de l'équipement ; et ledit moteur d'inférence (28) formant une interface avec l'opérateur au moyen de l'interface utilisateur et comprenant des programmes de diagnostic récursifs qui opèrent sur les données contenues dans la base de connaissances (32) pour générer des règles de diagnostic de symptômes en réponse à des symptômes introduits par l'opérateur qui sont représentatifs des anomalies présentes dans l'équipement, de telle façon que chaque règle de diagnostic de symptôme corrèle un symptôme avec son anomalie et soit représentative de l'anomalie ;

caractérisé en ce que :

ladite base de connaissances (32) représente la taxonomie ou le déterminant du système (figure 4:55) de l'équipement, ledit déterminant ayant des cellules de connaissances (56-70), une cellule de connaissances représentant une fonction définie (F) qu'un composant (substrat S(F)) de l'équipement exécute, une cellule de connaissances pouvant être connectée à une autre cellule de connaissances par l'intermédiaire d'une liaison descendante (58-60) si les deux cellules sont au même niveau, ou par un gestionnaire de niveau (56-58) si les deux cellules sont à des niveaux différents, ledit déterminant (55) étant partitionné en des éléments taxonomiques, chaque élément taxonomique se rapportant à un sous-ensemble de cellules de connaissances connectées par des liens, ladite base de connaissances comprenant les fichiers suivants :

- (1) un fichier de taxonomie contenant le nom de chaque cellule de connaissances, les liaisons descendantes de la cellule, et le nombre de gestionnaires de niveaux ;
- (2) un fichier de gestionnaires de niveaux contenant des informations de gestionnaires de niveaux indiquant si les cellules de connaissances sont des cellules composites pouvant être décomposées en des cellules de connaissances plus petites, indiquant ainsi si un composant de l'équipement peut être décomposé en une pluralité de composants plus petits ;
- (3) un fichier de frontières contenant une liste de cellules de connaissances qui sont également des cellules frontières dans chaque élément taxonomique ;
- (4) un fichier de listes de conducteurs contenant des informations concernant les conducteurs qui interconnectent les composants de l'équipement ;
- (5) un fichier de conditions de saut de niveau contenant des informations pour chaque élément taxonomique concernant une procédure de test permettant de déterminer si une cellule anormale existe à l'intérieur de l'élément taxonomique ;
- (6) un fichier de messages de test pour stocker les sorties de toutes les fonctions dans un élément taxonomique, qui sont générées lorsqu'un message de test représentant un stimulus est appliqué aux composants représentés par ledit élément taxonomique ;
- (7) un fichier de conditions de validité contenant des informations concernant les conditions devant être satisfaites par les composants de l'équipement pour un fonctionnement correct ;
- (8) un fichier de valeurs d'état de fonctionnement normal contenant des informations con-



cernant des sorties correctes d'un composant lorsque le stimulus lui est appliqué ;

dans lequel, pour localiser un composant anormal, le moteur d'inférence examine une séquence spécifique (séquence de transitions d'automate) de cellules de connaissances ou automates (un automate (X, A) est une cellule de connaissances A dans un élément taxonomique X), dans lequel le moteur d'inférence passe à un autre automate selon qu'une cellule examinée par le moteur d'inférence est une cellule frontière ou une cellule composite.

2. Système expert selon la revendication 1, caractérisé en outre en ce que ladite base de connaissances (32) comprend en outre des fichiers contenant des actions de maintenance tenant compte d'un ensemble complexe de facteurs rencontrés par expérience, lesquelles actions peuvent être appliquées audit équipement.
3. Système expert selon la revendication 2, caractérisé en outre en ce que chacune desdites cellules de connaissances est identifiée de façon unique par un automate représentant l'emplacement d'une cellule de connaissances à l'intérieur du déterminant, et ledit moteur d'inférence (28) comporte des moyens pour générer une séquence de transitions d'automates.
4. Système expert selon la revendication 2, caractérisé en outre en ce que ledit système expert (24) comprend en outre un tableau noir (34) couplé entre ladite base de connaissances (32) et ledit moteur d'inférence (28) et des moyens pour accéder aux données stockées dans la base de connaissances (32) et extraire des données pour construire une partie dudit déterminant sur ledit tableau noir pour que le moteur d'inférence (28) effectue son diagnostic.
5. Système expert selon la revendication 4, caractérisé en outre en ce que ledit moyen de construction d'une partie dudit déterminant sur ledit tableau noir (34) comprend un gestionnaire de tableau noir (30) et ledit moteur d'inférence (28) comprend un moyen pour appliquer audit gestionnaire de tableau noir la partie dudit modèle conceptuel devant être extraite.

55

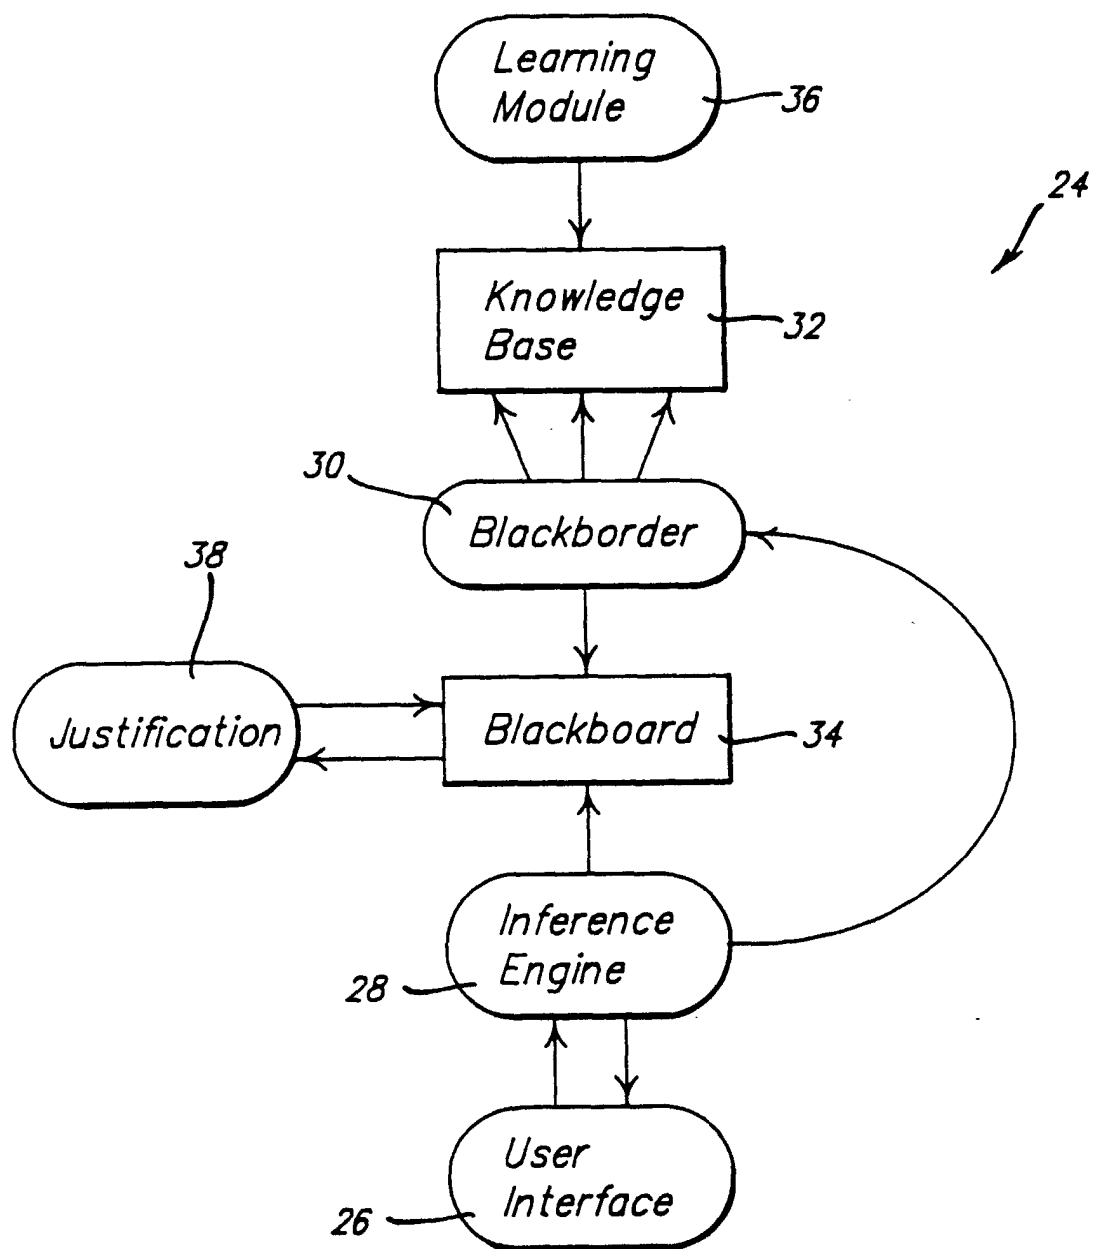


FIG. 1.




FIG. 2.





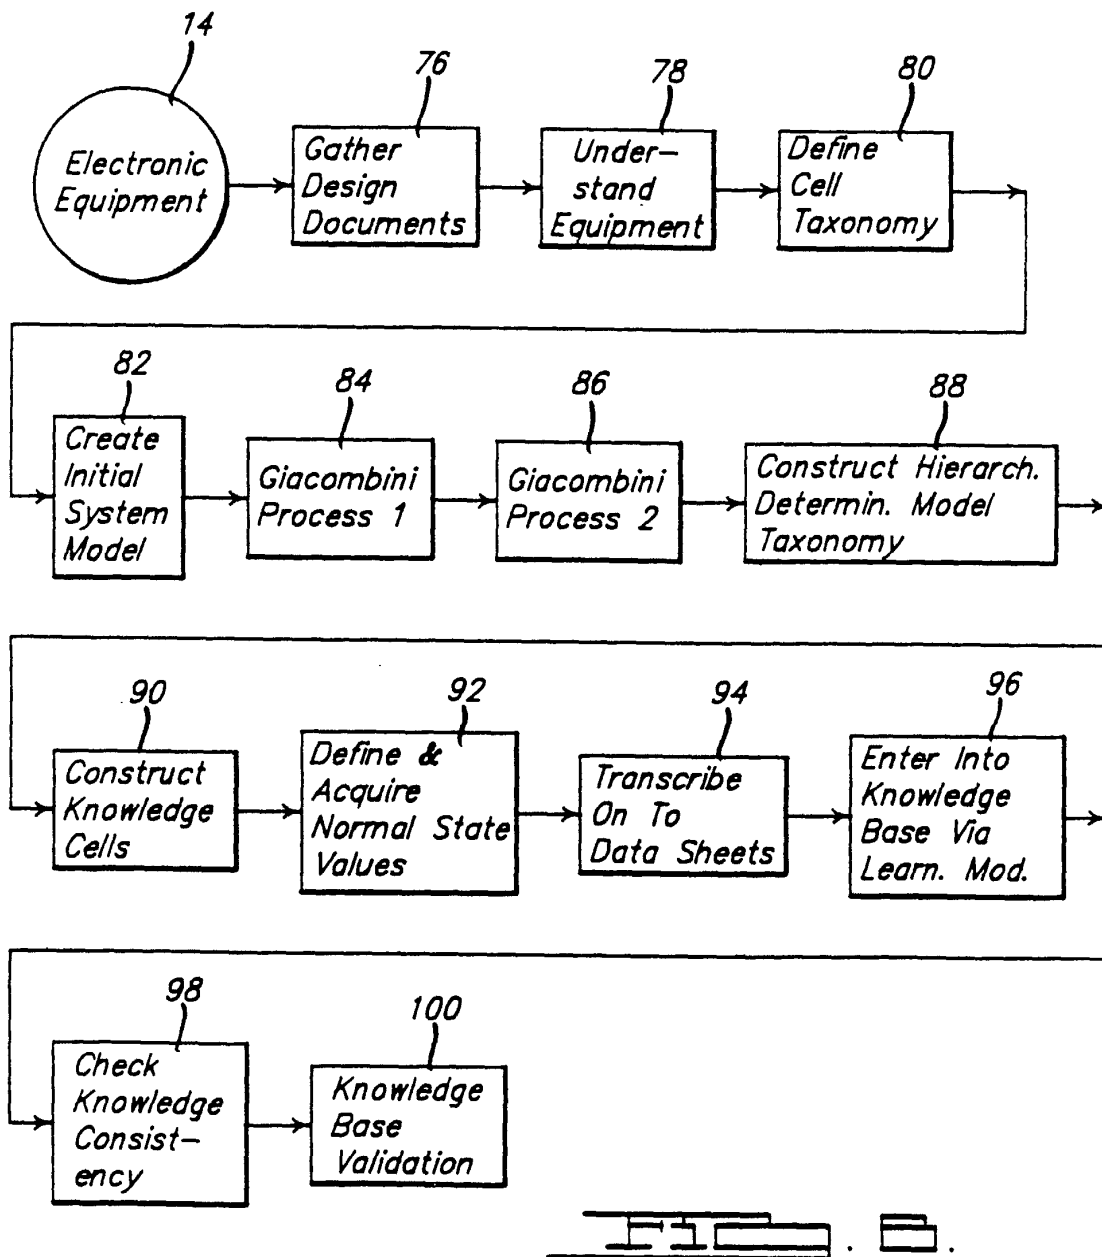



FIG. 3.








FIG. 2.






FIG. 3.






FIG. 11.




FIG. 12.






FIG. 13.

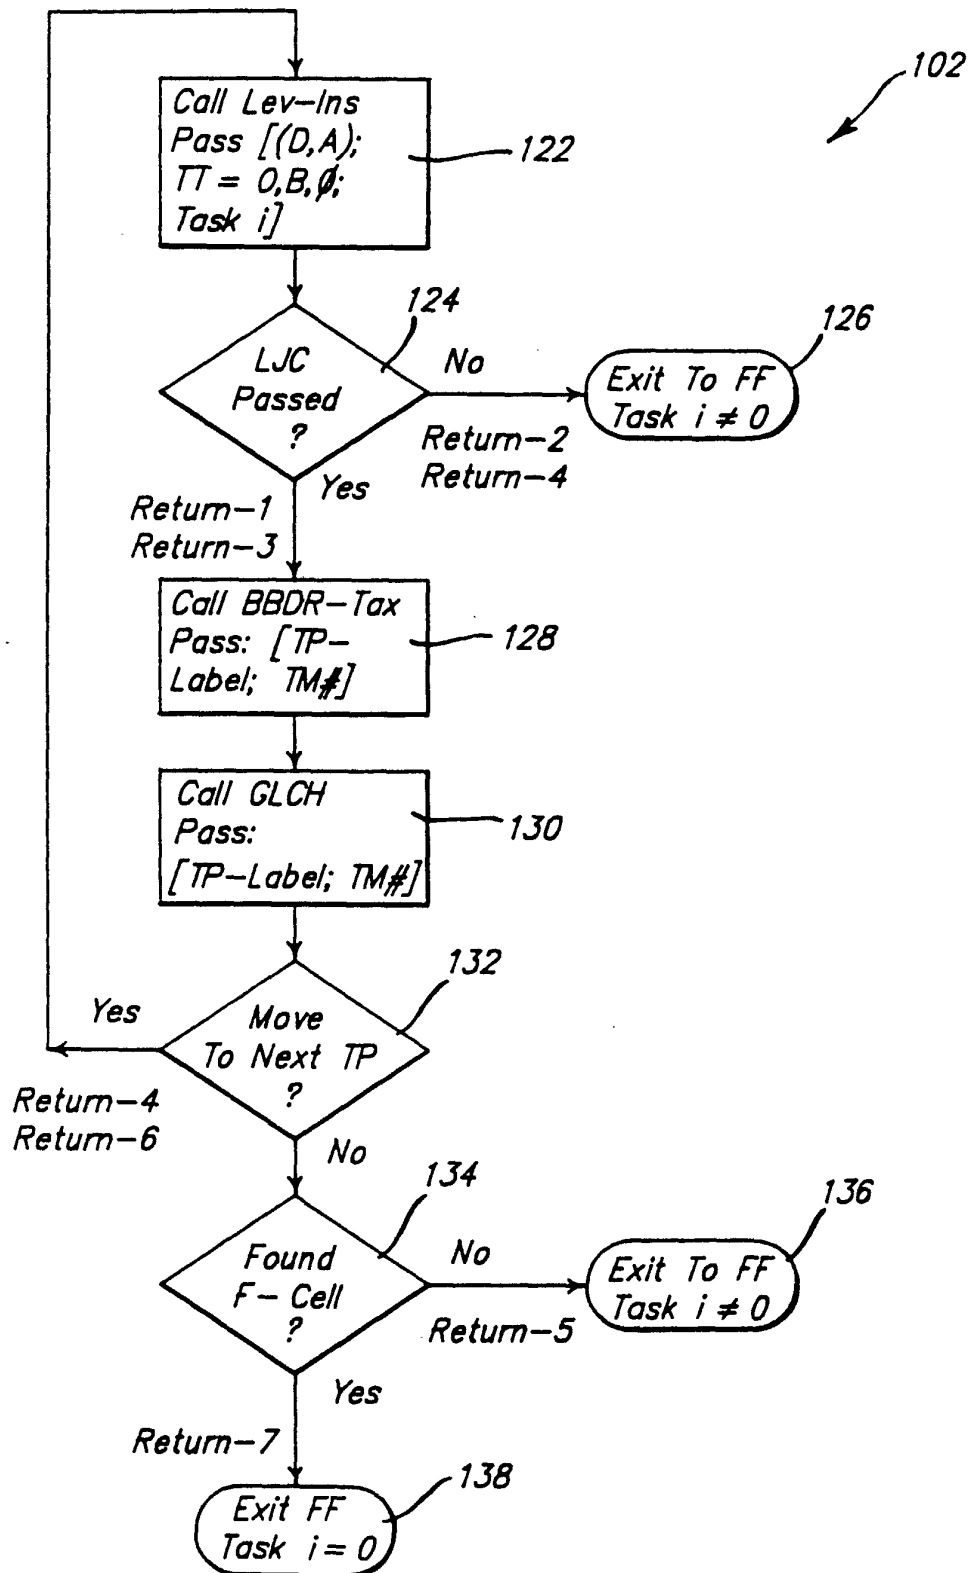


FIG. 14.

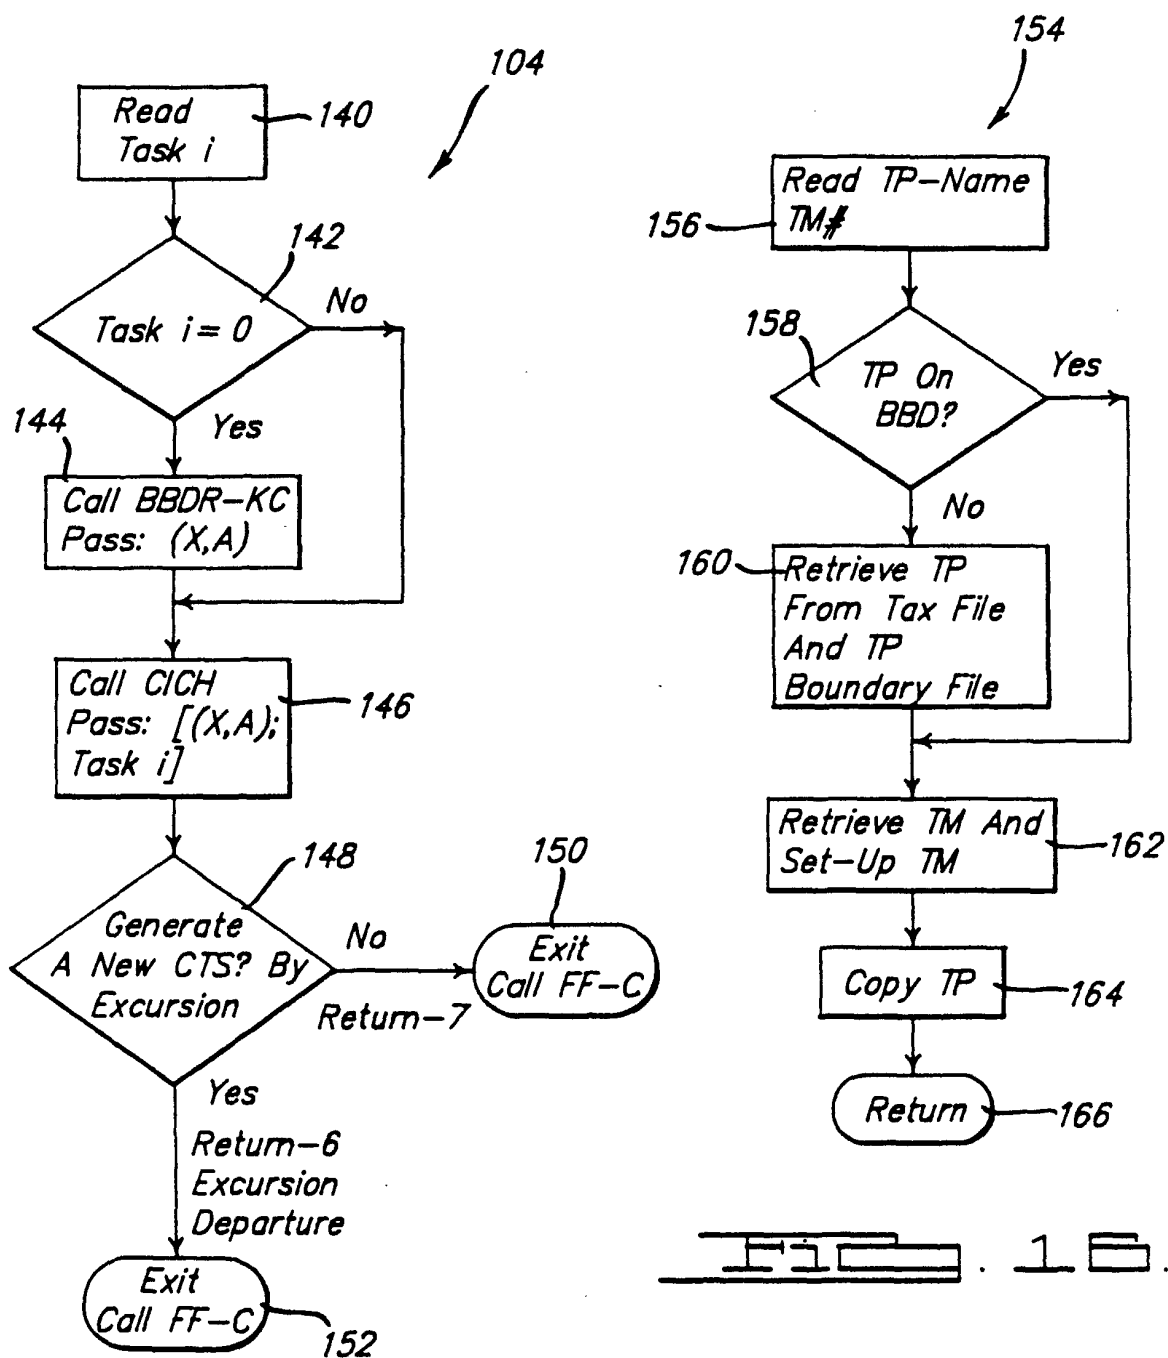


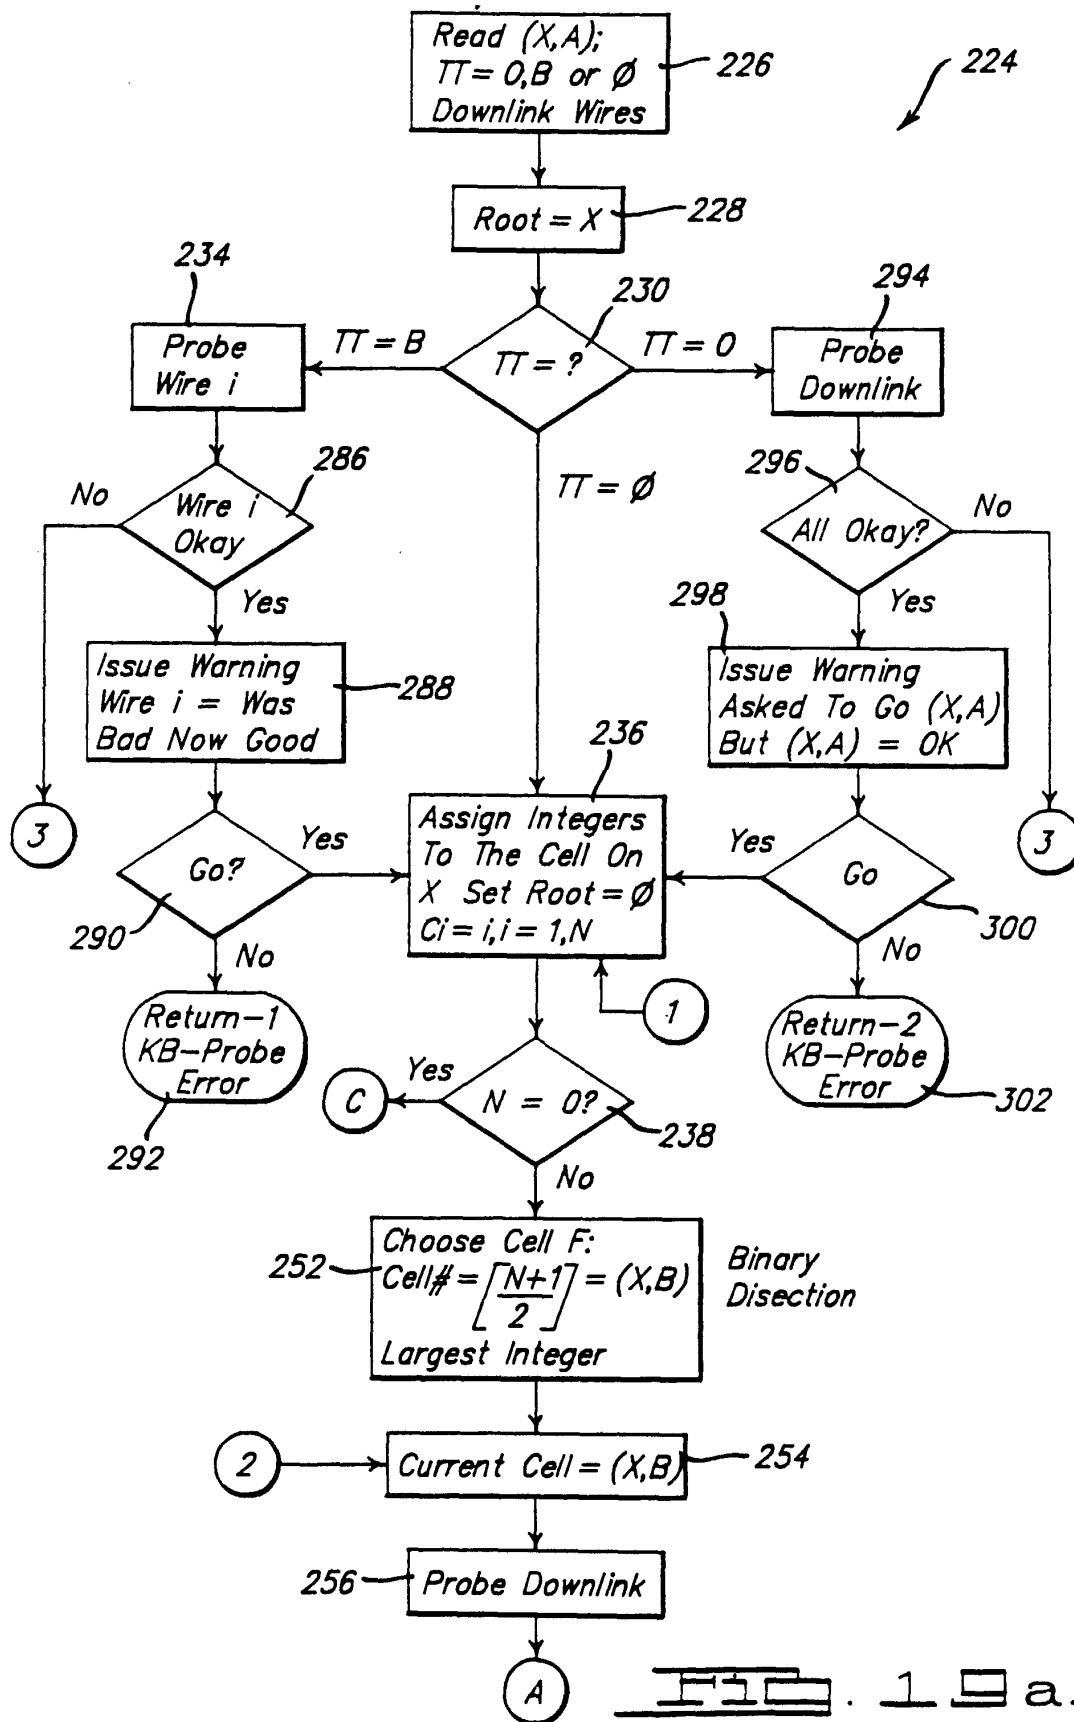



FIG. 15.

FIG. 16.

FIG. 17.

FIG. 18.









FIG. 19c.





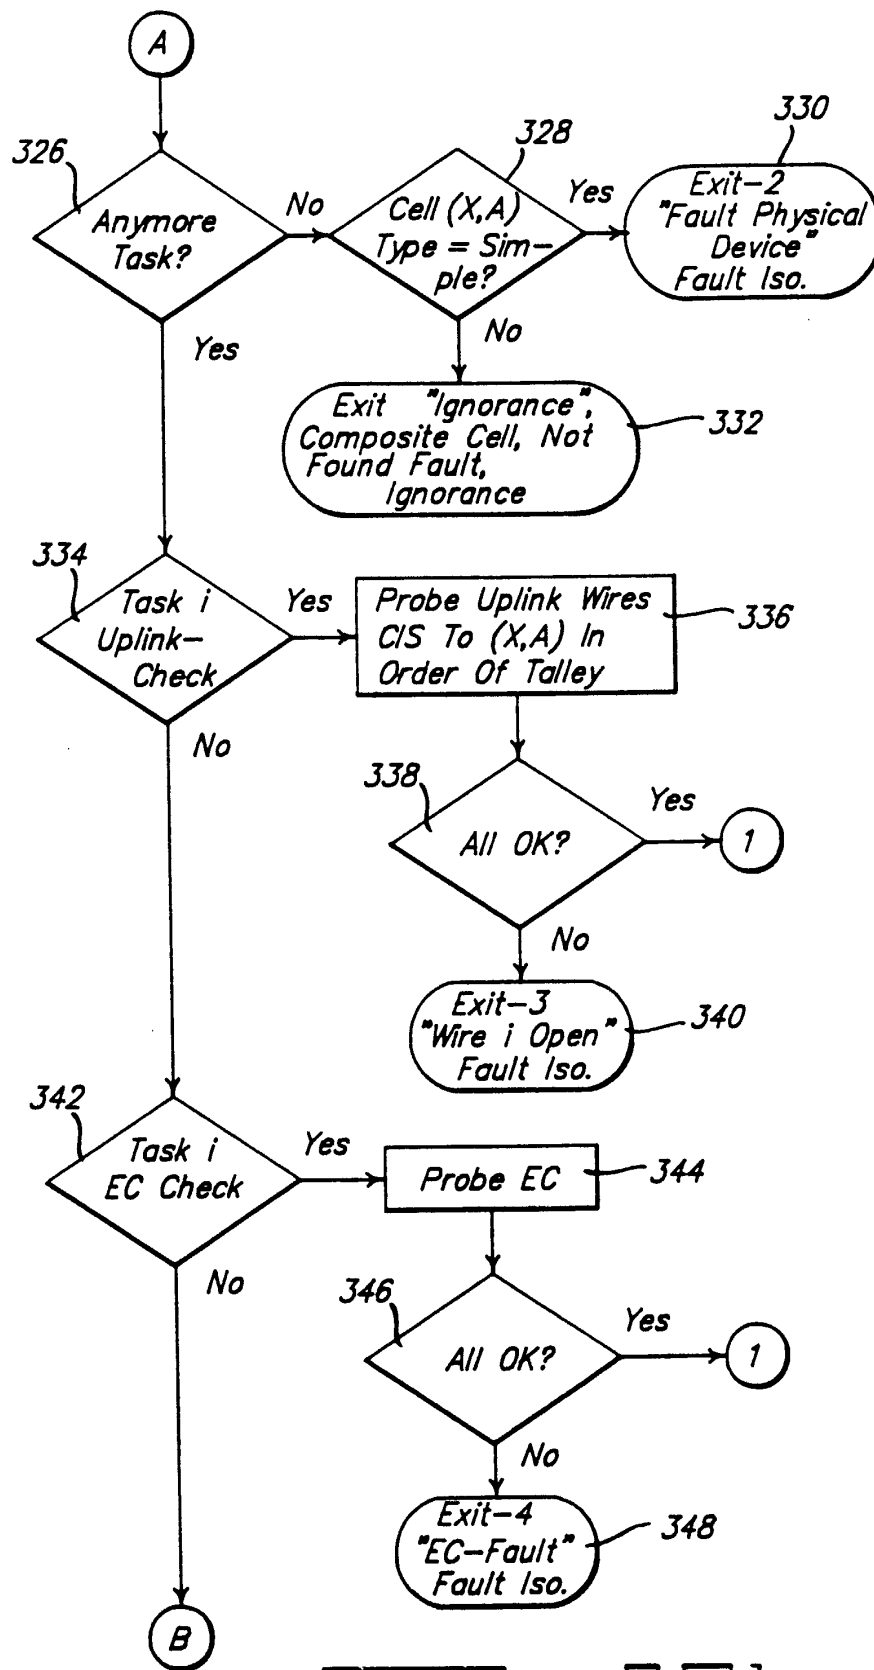


FIG. 2 . □ □ b.

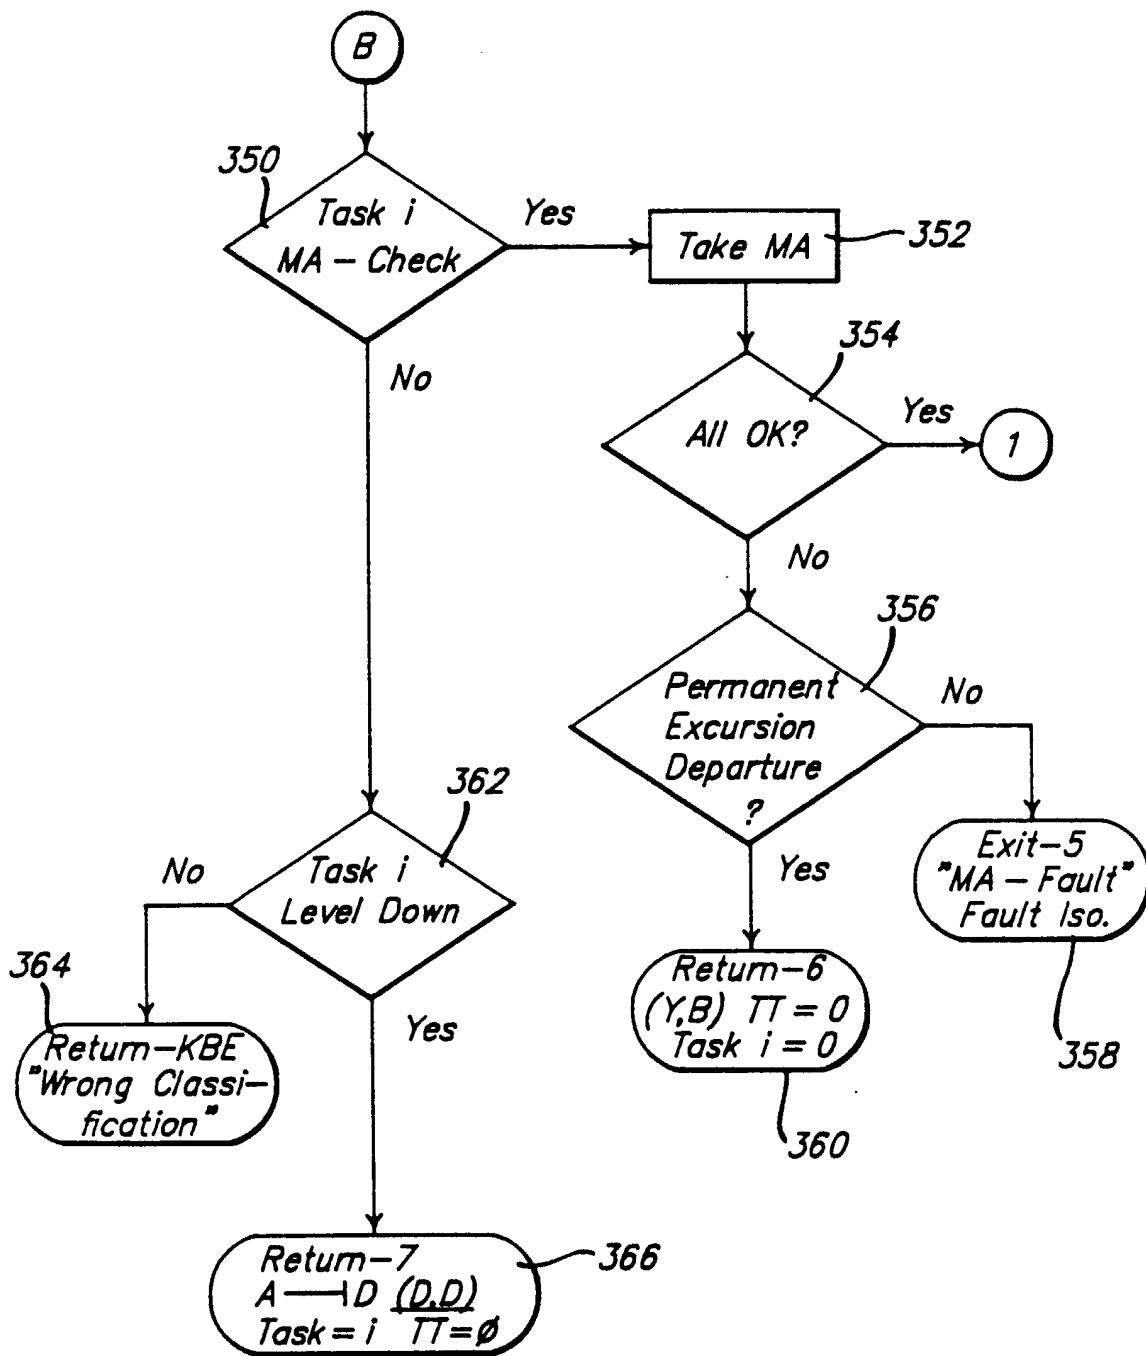


FIG. 20c.